Computational Materials Science 227 (2023) 112261

Contents lists available at ScienceDirect

Computational Materials Science

journal homepage: www.elsevier.com/locate/commatsci

Full length article

Convolutional neural networks for expediting the determination of minimum volume requirements for studies of microstructurally small cracks, part II: Model interpretation

Karen J. DeMille, Ashley D. Spear *

Department of Mechanical Engineering, University of Utah, Salt Lake City, 84112, UT, USA

ARTICLE INFO

Dataset link: https://doi.org/10.17632/nrntzs2zb7.1

Keywords:
Machine learning
Fracture mechanics
Representative volume element
Materials informatics
Saliency map

ABSTRACT

A significant question in the study of microstructurally small cracks (MSCs) is: *What is the minimum microstructural volume that should be included in studies involving MSCs?* To answer this, representative volume elements for microstructurally small cracks (RVE_{MSC}), or the minimum volume of microstructure required around an MSC to achieve convergence of crack-front parameters with respect to volume size, were previously determined using finite element (FE) simulations. The large computational expense of determining RVE_{MSC} via FE simulations motivated the implementation of convolutional neural networks (CNNs) to expedite the determination of RVE_{MSC} (Part I). In addition to expediting the determination of RVE_{MSC}, trained CNNs provide the opportunity to gain insights about RVE_{MSC} predictions through various interpretation methods, which we investigate in the current work. First, an inspection of CNN predictions reveals trends learned by the CNN. Second, an input sampling grid study offers insights into the volume of microstructure around an MSC that most influences predictions of RVE_{MSC}. Third, an input feature sensitivity analysis compares the influence of microstructural and geometrical features on RVE_{MSC} predictions. Fourth, visual inspections of saliency maps reveal the local microstructure that is most important to the CNN when predicting RVE_{MSC}. The CNN interpretation results show that microstructural features are more critical than geometrical features to the CNN predictions. Despite inherent limitations in interpreting saliency maps, the results demonstrate that the CNN can learn to identify various microstructural arrangements at individual crack-front points. Overall, this study highlights the importance of considering a variety of microstructural instantiations when determining RVE_{MSC}, as RVE_{MSC} should be a conservative minimum volume requirement that applies across a wide range of microstructural instantiations.

1. Introduction

Microstructurally small cracks (MSCs) are cracks whose lengths are on the order of the size of the predominant microstructural features [1]. Given the strong dependence of MSC behavior on microstructural features [2–14], there is a need to know how much volume of heterogeneous microstructure should be included in studies of MSCs. In a previous study, DeMille and Spear [15] addressed this issue by determining representative volume elements for microstructurally small cracks (RVE_{MSC}). By definition, RVE_{MSC} is "the smallest heterogeneous volume containing an MSC such that local crack-front parameters are converged with respect to volume size" [15]. Using a finite-element (FE) simulation framework, RVE_{MSC} was determined for linear-elastic microstructures, and it was found that microstructural heterogeneity, material constraint, and crack size influence the volume required for J-integral values to converge with respect to volume size.

Although the FE-based framework was successful at identifying RVE_{MSC} sizes and key influences on volume requirements, a significant limitation of the FE-based determination of RVE_{MSC} was the requisite computational expense; thus, a method for expediting the determination of RVE_{MSC} was proposed by DeMille and Spear in Part I [16]. In the proposed method, convolutional neural networks (CNNs) were trained to predict $RVE_{MSC,ip}$ sizes given microstructural and geometrical inputs, where $RVE_{MSC,ip}$ is the minimum volume required for the J-integral at a crack-front point p in microstructural instantiation i to converge with respect to volume size. The CNN predictions of $RVE_{MSC,ip}$ sizes were used to make CNN-based estimates of $RVE_{MSC,i}$ (the minimum volume required for the J-integral values at all crack-front points in microstructural instantiation i to converge with respect to volume size) by taking the maximum $RVE_{MSC,ip}$ size among all crack-front points in microstructural instantiation i. The rapidly determined CNN-based

* Corresponding author.
E-mail address: ashley.spear@utah.edu (A.D. Spear).

https://doi.org/10.1016/j.commatsci.2023.112261
Received 1 November 2022; Received in revised form 9 May 2023; Accepted 12 May 2023
Available online 3 June 2023

Nomenclature

a	Half-crack length
a/g	Normalized crack length
d_1	In-plane distance from a crack front to the sides of a volume
$d_{1,i}$	d_1 value corresponding to $\text{RVE}_{\text{MSC},i}$
$d_{1,ip}$	d_1 value corresponding to $\text{RVE}_{\text{MSC},ip}$
d_2	Distance from crack plane to the top and bottom of microstructural volume
$d_{2,i}$	d_2 value corresponding to $\text{RVE}_{\text{MSC},i}$
$d_{2,ip}$	d_2 value corresponding to $\text{RVE}_{\text{MSC},ip}$
d_{ahead}	Distance by which CNN sampling grid extends radially ahead of crack front
d_{behind}	Distance by which CNN sampling grid extends radially behind crack front
d_{cfp}	Euclidean distance from sampling grid point to crack-front point
d_{fs}	Distance from crack-front point to free surface in crack plane
Δ_{grid}	Grid spacing in CNN sampling grid
E	Elastic modulus
g	Average grain size
MSE	Mean squared error
N	Number of points per dimension in CNN sampling grid
R^2	Coefficient of determination
RVE_{MSC}	Representative volume element for a microstructurally small crack
$\text{RVE}_{\text{MSC},i}$	Minimum volume requirement for a microstructure i
$\text{RVE}_{\text{MSC},ip}$	Minimum volume requirement for crack-front point p in microstructure i
θ_{crit}	Position along crack front with the largest prediction of an $\text{RVE}_{\text{MSC},ip}$ parameter

estimates of $\text{RVE}_{\text{MSC},i}$ were then used to downselect microstructural instantiations predicted to have large $\text{RVE}_{\text{MSC},i}$ sizes. The downselected instantiations could then be simulated using FE analysis to determine RVE_{MSC}, where RVE_{MSC} ensures convergence of J-integral values at all crack-front points in all microstructural instantiations considered. Through the analysis completed in Part I [16], CNNs were shown to have the potential to expedite the computationally expensive, FE-based framework for determining RVE_{MSC}.

Despite key findings on CNN performance from Part I [16], the impacts of local microstructure and geometry features on CNN model predictions of volume requirements for J-integral convergence with respect to volume size were not revealed during the assessment of CNN performance. For example, the analysis of CNN performance did not answer questions such as: *What is the relative importance of microstructure and geometry in predicting* $RVE_{MSC,ip}$ *and* $RVE_{MSC,i}$*? What is the size of the region of local microstructure around a crack-front point that contains the grains that most influence CNN predictions of* $\text{RVE}_{\text{MSC},ip}$ and $\text{RVE}_{\text{MSC},i}$? *Which grains or arrangements of grains have the most influence on CNN predictions of* $RVE_{MSC,ip}$ *and* $RVE_{MSC,i}$*?*

Although CNNs have been demonstrated to be very adept at handling complex microstructures (e.g., to predict effective responses of microstructures [17–23], to segment microstructural images [24–29], to classify microstructures [30], to detect composite fibers [31], and to predict local material behavior [32–35]), CNNs are known as "black-box" machine learning (ML) models [26,36,37]. Namely, the CNN takes input and provides an output (or prediction) without directly informing the user why an output was predicted. However, methods for gaining insights into the predictions made by a CNN are available [36]. Beniwal et al. [18] and Eslamibidgoli et al. [38] used class activation mapping to visualize the regions of input microstructures that most contribute to CNN predictions. Similarly, Pokuri et al. [39] used saliency maps to visualize the regions of input microstructures that contributed the most to CNN predictions. Cecen et al. [17] visualized convolutional layer filter weights to identify microstructural patterns that a CNN searched for when making predictions. Kantzos et al. [40] visualized the output of intermediate CNN layers to identify regions of an input image that had the most influence on CNN predictions. Rahman et al. [41] used an occlusion sensitivity analysis, while Sung [42] used a sensitivity analysis, fuzzy curves, and a change of mean squared error (MSE) method to explore the importance of neural network inputs.

Part II of this work aims to explore the factors that influence the CNN predictions of $\text{RVE}_{\text{MSC},ip}$ and the CNN-based estimates of $\text{RVE}_{\text{MSC},i}$ made in Part I of this work [16]. The $\text{RVE}_{\text{MSC},i}$ and $\text{RVE}_{\text{MSC},ip}$ data used in this work were obtained from the original RVE_{MSC} study by the authors [15]. The following sections discuss various interpretations of the CNNs implemented in Part I [16]. First, predictions from five-fold cross-validations are inspected to discover general trends in CNN predictions and estimates of $\text{RVE}_{\text{MSC},ip}$ and $\text{RVE}_{\text{MSC},i}$ parameter values. Second, input sensitivity analyses are conducted to provide a model-level interpretation of the relative importance of different inputs in predicting $\text{RVE}_{\text{MSC},ip}$ parameter values. Third, an input sampling grid study provides a model-level interpretation of the region surrounding a crack-front point that contains the grains that most influence the CNN predictions of $\text{RVE}_{\text{MSC},ip}$ parameter values. Fourth, saliency maps are generated to provide visualizations of the relative importance of local microstructural arrangements in predicting $\text{RVE}_{\text{MSC},ip}$ parameter values.

2. Methods

2.1. Previous work: RVE_{MSC} data

The data [43] used in Part I [16] and again in this work were obtained from a previous study by the authors [15] and are briefly described here for completeness. In the previous study, FE models of cracked microstructures were simulated to determine RVE_{MSC} for linear-elastic, heterogeneous domains. Each idealized microstructural volume contained a static, semi-circular, planar surface crack of half-length a, surrounded by cube-shaped grains with side length g and elastic moduli E ranging from 75 to 225 GPa. The volume of microstructure surrounding a given crack was parameterized using two values: d_1 (the minimum in-plane distance between the crack front and the sides of the volume) and d_2 (the distance between the plane of the crack and the upper and lower surfaces of the volume). For generalization, half-crack lengths and volume parameters were expressed relative to grain size, g. The FE simulations provided J-integral values[1] that were used to determine $\text{RVE}_{\text{MSC},ip}$, $\text{RVE}_{\text{MSC},i}$, and RVE_{MSC} sizes under two different types of boundary conditions. Key aspects, assumptions, and relevant outputs of the FE simulations from the original RVE_{MSC} study [15], including idealized microstructure representation, volume parameterization, and boundary condition sets, are shown in Fig. 1. For complete details of the models, simulations, calculation of volume requirements, and analysis, see DeMille and Spear [15] and Part I [16] of this work.

[1] J-integral values corresponding to nodes located at grain boundaries were not considered in the FE-based determination of volume requirements due to spurious J-integral values that occur at grain boundaries.

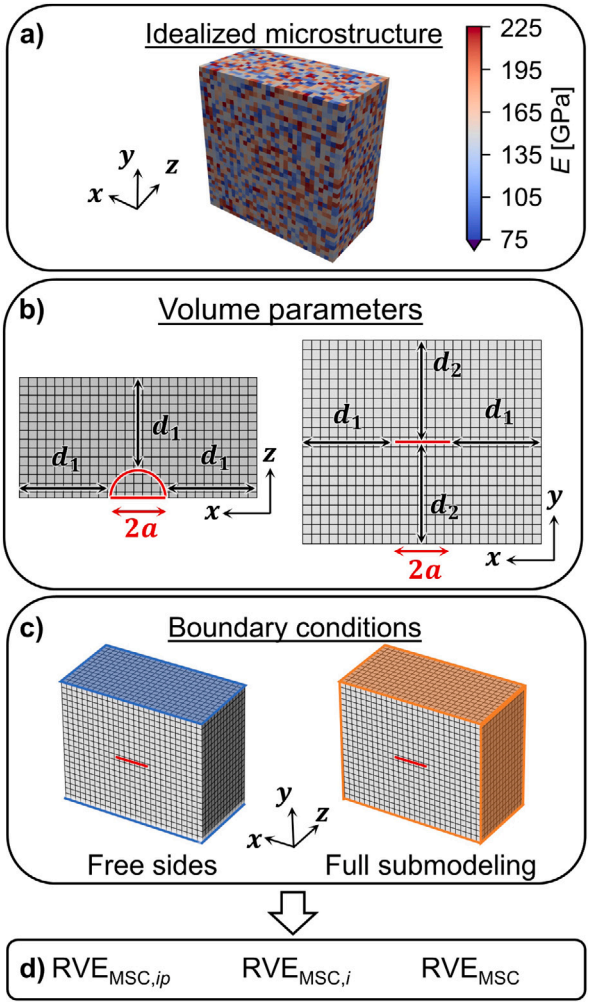

Fig. 1. Key features of FE simulations used to generate RVE_{MSC} data during a previous study [15]: (a) Idealized, heterogeneous microstructures consisted of cube-shaped grains of side length g with elastic moduli values ranging from 75 to 225 GPa. (b) Two parameters (d_1 and d_2) were used to describe the volume of microstructure around a crack. (c) Two types of boundary conditions were investigated; displacements applied to the highlighted surfaces in each case were interpolated from a global model under uniaxial tension. (d) Results from FE simulations were used to identify volume requirements for J-integral convergence at a given crack-front point in a given microstructure ($RVE_{MSC,ip}$), for an entire given microstructure ($RVE_{MSC,i}$), and for an entire population of given microstructures (RVE_{MSC}).

2.2. Convolutional neural network

In Part I [16], CNNs [44,45] were implemented to expedite the determination of RVE_{MSC} using the data described above. In the current work, the same CNNs are interpreted to explore the mechanisms governing the predictions made regarding J-integral convergence with respect to volume size. Several of the interpretation methods presented in this work are inherently tied to the model training process; therefore, rather than simply using the trained models or results from Part I, the CNN models are re-trained in this study for the purpose of model interpretation. For completeness, the CNN model setups described in Part I, including inputs, architecture, and targets, are described again here but in less detail. Four different CNNs are used to predict two $RVE_{MSC,ip}$ parameters ($d_{1,ip}$ and $d_{2,ip}$) for each of the two boundary condition types (free sides and full submodeling). The reader is referred to Part I [16] for a more detailed description.

2.2.1. CNN inputs

Fig. 2a shows the four features that define the microstructural inputs (viz., local elastic modulus values (E)) and geometrical inputs (viz., normalized crack size (a/g), distance to the free surface (d_{fs}), and distance to crack-front point (d_{cfp})) for a given crack-front point p in microstructural instantiation i. The inclusion of the E and d_{fs} inputs reflect observations that scatter in small crack behavior is caused by variations in microstructure and that small crack behavior is strongly dependent on free surfaces, respectively [2,7,46]. The inputs a/g and d_{fs} are scalars, where d_{fs} is the distance in the z-direction between p and the unconstrained model face at $z = 0$. The inputs E and d_{cfp} are 3D arrays of microstructural or geometrical features sampled from a grid surrounding crack-front point p. Elastic modulus values are sampled for the E array, while Euclidean distances between a grid point and crack-front point p are sampled for the d_{cfp} array.

The sampling strategy for obtaining the E and d_{cfp} arrays involves $N \times N \times N$ grids located at each crack-front point p and oriented tangent to the crack front. Various sampling grids are shown in Fig. 3. The coverage of a sampling grid is defined through three parameters: d_{ahead}, d_{behind}, and N. As shown in Fig. 3, d_{ahead} is the distance the grid extends ahead of the crack-front point, d_{behind} is the distance the grid extends behind the crack-front point, and N is the number of grid points in each direction. The grid spacing, Δ_{grid}, in each direction can be derived from the previous three parameters using the equation:

$$\Delta_{grid} = \frac{d_{ahead} + d_{behind}}{N - 1}. \quad (1)$$

Unless otherwise noted, each sampling grid for this work contains 16 × 16 × 16 points and extends 4.00 grains ahead of and 0.95 grains behind a given crack-front point. Thus, the grid spacing is 0.33 grains. For more details on the selection of these grid parameters, see Part I [16]. With this sampling strategy, some sampling points fall outside of the FE model volume, in which case a value of −1 is assigned to those points in the input arrays (shown in purple in Fig. 2a). Before being input to the CNN, both scalar inputs and array inputs for grid points that fall within the FE model volume are normalized to a range of 0 to 1.

2.2.2. CNN architecture

The CNN architecture consists of three blocks of convolutional and pooling layers followed by dropout and fully connected layers, as shown in Fig. 2b. The 3D arrays are input to the first block of convolutional/pooling layers. The scalar inputs are concatenated with the flattened output from the third block of convolutional/pooling layers prior to the first fully connected layer. The predicted $RVE_{MSC,ip}$ parameter value is output from the second fully connected layer. The blocks of two convolutional layers followed by a maximum pooling layer are based on VGG-16 [47]. Each CNN is trained using the MSE as the loss function. The CNNs are implemented using Keras [48] and TensorFlow [49]. For more details of the CNN architecture and training parameters, see Part I [16].

2.2.3. CNN targets

As shown in Fig. 2c, the targets of the CNN models are $RVE_{MSC,ip}$ parameter values ($d_{1,ip}$ and $d_{2,ip}$). Four separate CNNs are trained, where each CNN predicts either $d_{1,ip}$ or $d_{2,ip}$ values under either free-sides or full-submodeling boundary conditions. Using four separate CNNs is similar to the approach used by Rovinelli et al. [50,51], where separate ML models were used to predict different crack-growth parameters. The target $RVE_{MSC,ip}$ parameters are calculated from the FE simulations completed in the previous RVE_{MSC} determination study by the authors [15]. The calculation of $RVE_{MSC,ip}$ parameter values is described in full detail in Part I [16]. The target $RVE_{MSC,ip}$ parameter values are normalized to a range of 0 to 1 for use in the CNN.

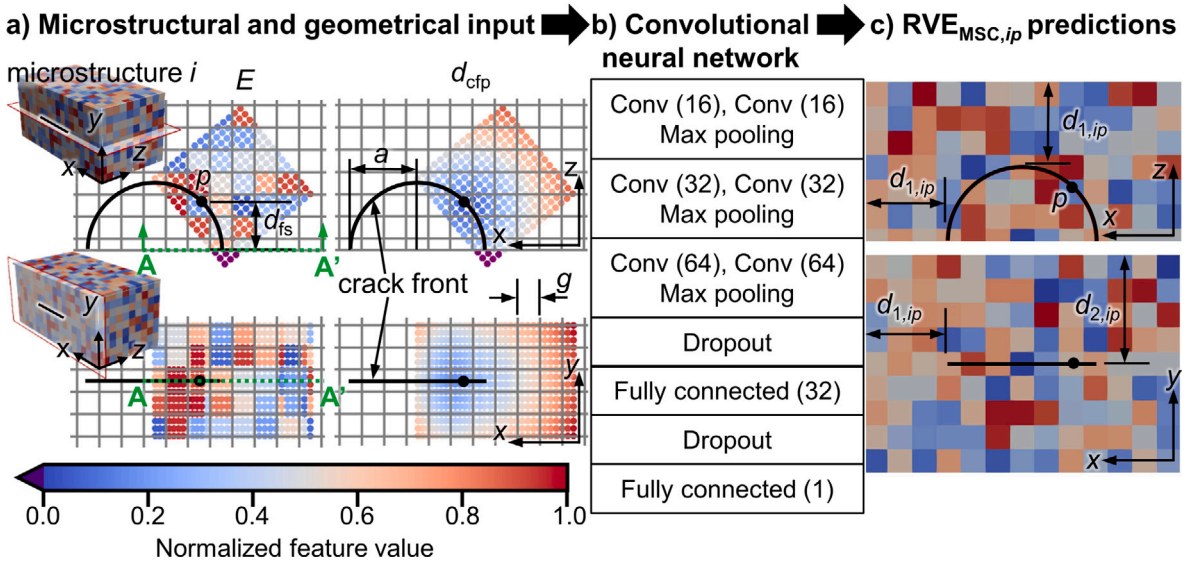

Fig. 2. Summary of the CNNs used in Part I [16] and re-trained in this work using various methods for model interpretation. The main components of the CNNs are: (a) microstructural and geometrical input features, (b) CNN architecture, and (c) RVE$_{MSC,ip}$ parameter predictions. Input and output features are obtained from FE simulations completed during the original RVE$_{MSC}$ study [15].

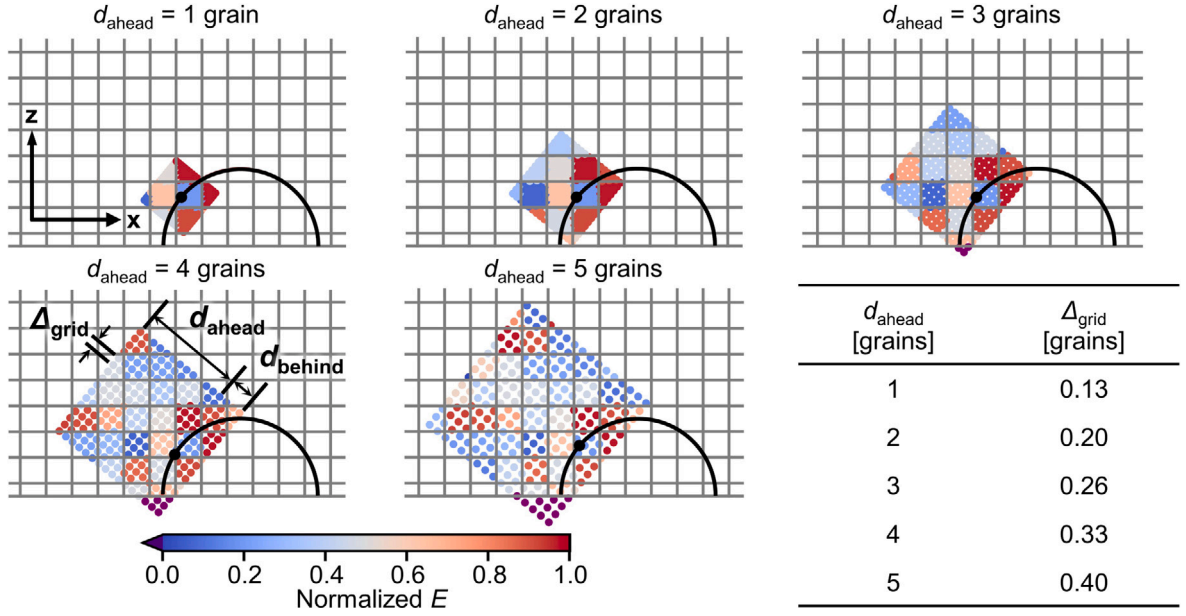

d_{ahead} [grains]	Δ_{grid} [grains]
1	0.13
2	0.20
3	0.26
4	0.33
5	0.40

Fig. 3. The five sampling grids used to sample 3D array inputs during the input sampling grid study. In the study, the distance sampled ahead of the crack front (d_{ahead}) is varied from 1 to 5 grains. The resulting grid point spacing for each sampling grid is shown on the bottom right.

2.2.4. CNN-based estimates of RVE$_{MSC,i}$ and RVE$_{MSC}$

The CNN-based predictions of RVE$_{MSC,ip}$ at individual crack-front points are used to establish minimum volume requirements for individual microstructural instantiations (RVE$_{MSC,i}$) and for all considered microstructural instantiations (RVE$_{MSC}$) by taking, respectively, the maximum RVE$_{MSC,ip}$ parameter values across all crack-front points in microstructural instantiation i:

$$d_{1,i} = \max_{p}(d_{1,ip}) \tag{2}$$

$$d_{2,i} = \max_{p}(d_{2,ip}) \tag{3}$$

and subsequently the maximum RVE$_{MSC,i}$ parameter values across all microstructural instantiations:

$$d_{1,MSC} = \max_{i}(d_{1,i}) \tag{4}$$

$$d_{2,MSC} = \max_{i}(d_{2,i}). \tag{5}$$

2.2.5. CNN data overview and splitting

In total, 88 cracked microstructural instantiations were simulated using FE analysis during the previous study [15]: 22 microstructural instantiations for each of four different normalized crack sizes (a/g = 0.25, 0.45, 1.0, 3.0). Depending on the size of the crack, each cracked microstructural instantiation contained between 91 and 269 crack-front points (excluding points located at grain boundaries). Each set of crack-front points from a given cracked microstructural instantiation constitutes one crack-front point (CFP) group. Each of the four CNNs (one for each combination of boundary condition and RVE$_{MSC,ip}$ parameter) has a total of 88 CFP groups of data points available for CNN training and testing. Altogether, these 88 CFP groups contain a total of

Table 1
Input feature combinations considered during an input feature sensitivity analysis. During the sensitivity analysis, a repeated five-fold CV is performed for each of the input feature combinations.

Dropped input	Array Inputs	Scalar Inputs
None	E, d_{cfp}	a/g, d_{fs}
E	d_{cfp}	a/g, d_{fs}
d_{cfp}	E	a/g, d_{fs}
a/g	E, d_{cfp}	d_{fs}
d_{fs}	E, d_{cfp}	a/g

14 168 data points. When splitting the data for training and testing, the data are split according to the CFP group (i.e., points within each CFP group are kept together). This splitting approach ensures that the CNN remains unbiased during testing, as similar data points from a given cracked microstructure are not contained in both the training and test set.

2.2.6. Five-fold cross-validation

A five-fold cross-validation (CV) is used to assess the performance of the CNN models. The 88 CFP groups are split into five folds containing 17, 17, 17, 17, and 20 CFP groups, respectively. The model is then trained and tested using three folds for training, one fold for validation, and the remaining fold for testing. This training and testing process is repeated five times, using each fold as the test set during one of five independent training runs. At the beginning of each independent training run, model weights are re-initialized. Each CFP group is used as test data exactly once during the CV. In other words, the five-fold CV provides five estimates of model performance (one for each independent training run) and a blind prediction of an $\mathrm{RVE}_{\mathrm{MSC},ip}$ parameter for each crack-front point, where each blind prediction comes from one of the five models trained during the CV. These blind predictions of $\mathrm{RVE}_{\mathrm{MSC},ip}$ parameters can then be used to determine blind estimates of $\mathrm{RVE}_{\mathrm{MSC},i}$ parameters for each cracked microstructural instantiation.

2.3. CNN interpretation studies

2.3.1. Input feature sensitivity analysis

A sensitivity analysis is performed to determine which of the four CNN input features most influences the prediction of $\mathrm{RVE}_{\mathrm{MSC},ip}$ values. In this study, the CNNs are trained using five different combinations of input features. The first combination uses all four of the input features. In the remaining input combinations, one of the four input features is omitted from the CNN. The five input feature combinations are shown in Table 1. The approach used for the sensitivity analysis is similar to an occlusion sensitivity study performed by Rahman et al. [41] and a change of MSE study performed by Sung [42].

Five-fold CVs (Section 2.2.6) are used to assess the performance of the CNNs with various input features. For each input combination, the five-fold CV process is repeated five different times, using a different split of the 88 CFP groups each time. The CV is repeated to obtain more performance metrics over which to average and, thus, to better gauge model performance for each set of input parameters. Given that each CV provides five estimates of model performance, the CVs provide 25 different estimates of model performance for each input combination. The average of 25 model performance metrics is used to compare input combinations.

2.3.2. Input sampling grid study

The physical size of the 3D sampling grids (Section 2.2.1) is varied to compare the influence that microstructural features at various distances from the crack front have on volume requirement predictions. Five different sampling grids are considered, where each grid extends a different distance ahead of the crack front ($d_{\mathrm{ahead}} = 1, 2, 3, 4, 5$ grains). The grid parameters d_{behind} and N are kept fixed at 0.95 grains and 16, respectively. Fig. 3 shows the five sampling grids along with the grid spacing (Δ_{grid}) for each grid.

As in the input feature sensitivity analysis (Section 2.3.1), five-fold CVs (Section 2.2.6) are used to assess the performance of the CNNs with different input sampling grids. Repeated CVs are used to increase the number of available CNN performance estimates. For each sampling grid, five different splits of the 88 CFP groups are used in five independent five-fold CVs. The CVs provide 25 different estimates of model performance metrics for each sampling grid (one CV provides five performance estimates). The 25 performance metrics are averaged to compare CNN performance with inputs sampled using various d_{ahead} values.

2.3.3. Saliency maps

Saliency maps are extracted from trained CNN models to identify microstructural input voxels that are most critical to the prediction of $\mathrm{RVE}_{\mathrm{MSC},ip}$. The trained CNN models come from a five-fold CV (Section 2.2.6), where the CNN model used to generate a given saliency map is the model trained while the crack-front point of interest is held out as part of the test data. As described by Simonyan et al. [52], saliency maps are found by taking the absolute value of the derivative of a CNN model output with respect to the model input arrays via back-propagation. Given that the CNN models have multiple 3D input arrays, the maximum absolute value of the derivative over all input arrays is taken at each voxel location. The resulting saliency maps are 3D arrays with the same dimensions as the input arrays ($16 \times 16 \times 16$ for this work). The saliency maps are extracted using the package tf-keras-vis [53] alongside Keras [48] with a TensorFlow backend [49]. SmoothGrad [54] is used to reduce noise in the extracted saliency maps. For each CNN model, one saliency map is extracted at each crack-front point p in each microstructural instantiation i.

3. Results

3.1. Five-fold CV predictions/estimates

Fig. 4 shows actual and CNN-estimated $\mathrm{RVE}_{\mathrm{MSC},i}$ parameter values (assuming a 5 % J-integral convergence tolerance) for the 22 microstructural instantiations associated with each crack size. The actual $\mathrm{RVE}_{\mathrm{MSC},i}$ parameter values, shown in blue in Fig. 4, were found using the FE simulations described in previous work [15]. Meanwhile, CNN-estimated $\mathrm{RVE}_{\mathrm{MSC},i}$ parameter values, shown in red in Fig. 4, are obtained from blind $\mathrm{RVE}_{\mathrm{MSC},ip}$ predictions made during a five-fold CV. The sizes of the points reflect the number of times a given parameter value was repeated in the data. Results are shown for four normalized crack sizes ($a/g = 0.25,\ 0.45,\ 1.0,\ 3.0$), two boundary conditions (full submodeling and free sides), and two $\mathrm{RVE}_{\mathrm{MSC},i}$ parameters ($d_{1,i}$ and $d_{2,i}$). The results in Fig. 4 are used to evaluate the trends present in CNN-based $\mathrm{RVE}_{\mathrm{MSC},i}$ estimates, discussed in Sections 4.1, 4.2, and 4.7.

Fig. 5 shows actual and CNN-estimated $\mathrm{RVE}_{\mathrm{MSC},i}$ parameter values from a five-fold CV. Results are plotted for three types of microstructural instantiations. Parameter values for uniform and "2nn max/min E" microstructural instantiations are shown in red and cyan, respectively. Parameter values for all instantiations with varied E values are shown in blue, where the size of each point reflects the number of times a given parameter value is repeated. For uniform microstructural instantiations, all grains have $E = 138$ GPa. For 2nn max/min E instantiations, one cracked grain is assigned the maximum (or minimum) possible E value, and all of its second-nearest-neighbor grains are assigned the minimum (or maximum) possible E value. The second-nearest-neighbor grains are all grains that share one face with a first-nearest-neighbor grain. Similarly, first-nearest-neighbor grains share one face with a cracked grain. Note that the 2nn max/min E microstructural instantiations are also included in the set of varied E instantiations. The results in Fig. 5 are used to evaluate the

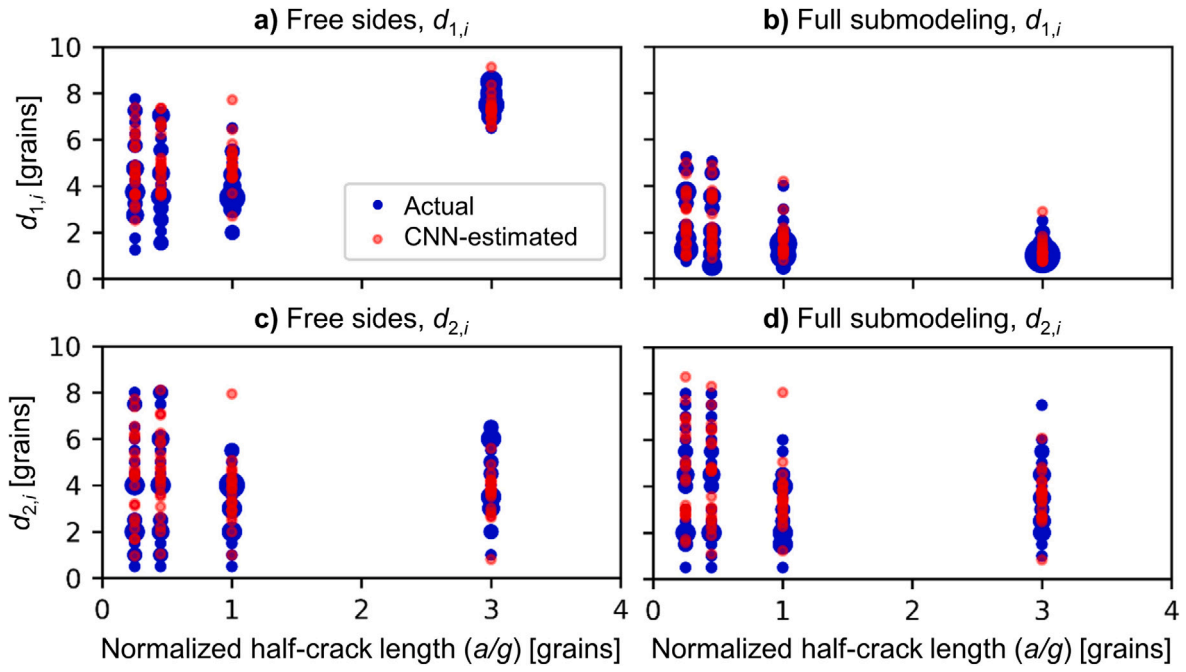

Fig. 4. Comparison between actual and CNN-estimated $\text{RVE}_{\text{MSC},i}$ parameter values ($d_{1,i}$ and $d_{2,i}$) for 22 microstructural instantiations per crack size under two sets of boundary conditions (free sides and full submodeling). The actual values were calculated in a previous RVE_{MSC} study by the authors [15]. The CNN-estimated $\text{RVE}_{\text{MSC},i}$ parameter values are obtained from CNN predictions of $\text{RVE}_{\text{MSC},ip}$ parameters made during a five-fold CV. The size of each point reflects the number of times a value is repeated in the set.

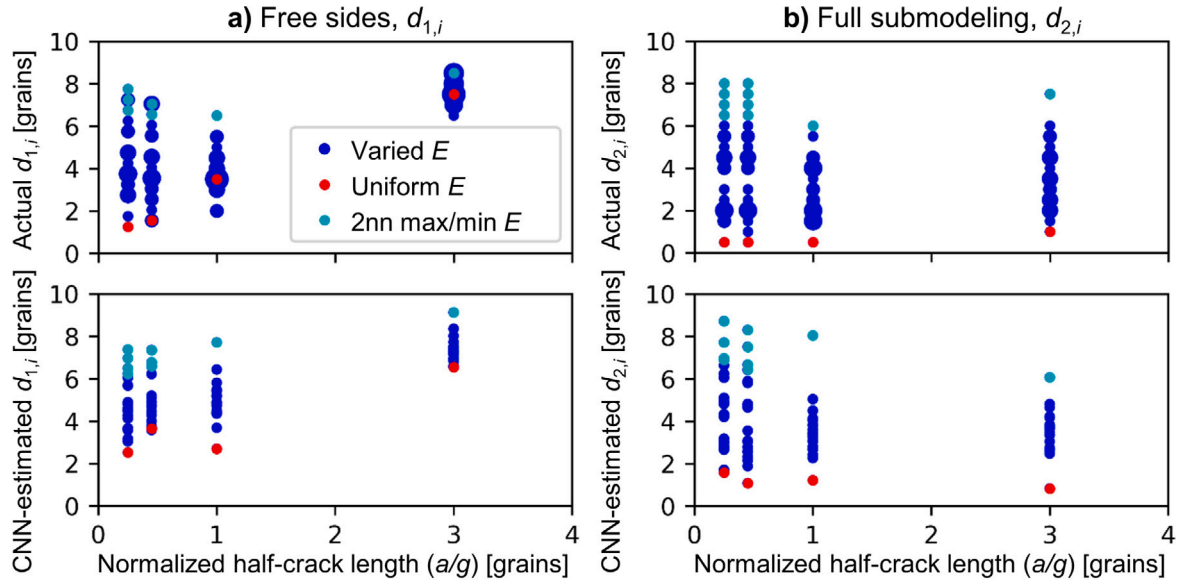

Fig. 5. Comparison between actual and CNN-estimated $\text{RVE}_{\text{MSC},i}$ parameter values for two sets of $\text{RVE}_{\text{MSC},i}$ parameters. Parameters shown are: (a) Free sides, $d_{1,i}$ and (b) Full submodeling, $d_{2,i}$. Parameter values for three specific types of microstructural instantiations are highlighted. Uniform and varied E instantiations have the same and different E value(s) assigned to grains, respectively. The 2nn max/min E instantiations have all second-nearest-neighbor grains of a one cracked grain assigned either the maximum or minimum E value. Actual parameter values were obtained from a previous study [15] and CNN-estimated parameter values are obtained from CNN predictions made during a five-fold CV. The size of each point represents the number of times a given value is repeated in the data set.

trends present in CNN-based $\text{RVE}_{\text{MSC},i}$ estimates of extreme volume requirements (discussed in Sections 4.4 and 4.8).

Fig. 6a compares the actual and predicted $d_{1,ip}$ values for the microstructural instantiation having the maximum CNN-estimated $d_{1,i}$ value for $a/g = 1$ under full-submodeling boundary conditions. In the plot, actual and predicted $d_{1,ip}$ values are plotted for each point along the crack front. A slice of the local microstructure (E array) in the neighborhood of the crack-front point with the maximum $d_{1,ip}$ prediction (corresponding to θ_{crit}) is shown in Fig. 6b. Recall, the $d_{1,ip}$ prediction that governs the CNN estimate of $d_{1,i}$ is the maximum $d_{1,ip}$ prediction over all crack-front points p in instantiation i. The dashed and solid green lines on the slice of microstructure serve as indicators for the orientation of the slice relative to the crack front. The results in Fig. 6 are used to evaluate the trends present in $\text{RVE}_{\text{MSC},ip}$ predictions for one illustrative microstructure (discussed in Sections 4.1 and 4.7) and to explore microstructural arrangements that influence CNN-based $\text{RVE}_{\text{MSC},i}$ estimates of extreme volume requirements (discussed in Section 4.4).

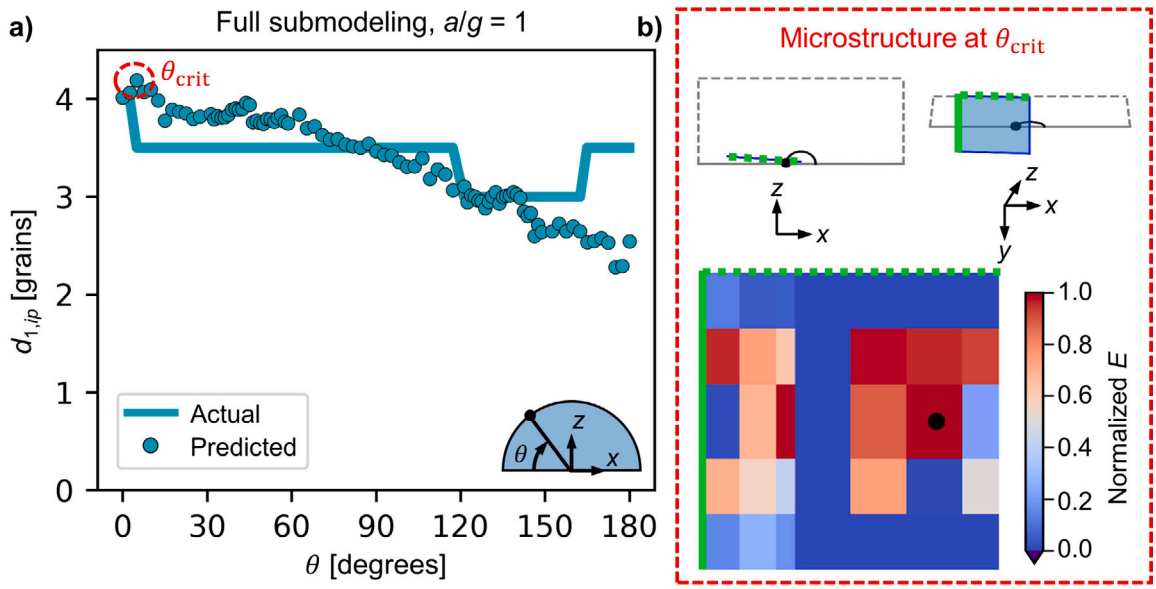

Fig. 6. A comparison between the actual and predicted $d_{1,ip}$ values for the microstructure having the maximum $d_{1,i}$ value when subjected to full-submodeling boundary conditions and containing a crack of size $a/g = 1$. (a) Actual and predicted $d_{1,ip}$ values are plotted as a function of position, θ, along the crack front. Actual values are from FE simulations [15], and predicted values are from a trained CNN model. The crack-front point corresponding to the largest predicted $d_{1,ip}$ value is denoted θ_{crit}. (b) The microstructural configuration in the neighborhood of θ_{crit} is depicted by a slice through the 3D microstructure; the position and orientation of the slice are shown at top, and the normalized E values around θ_{crit} are shown at bottom.

3.2. Input feature sensitivity analysis

The results of the input feature sensitivity analysis described in Section 2.3.1 are shown in Fig. 7. For each combination of boundary condition type and $RVE_{MSC,ip}$ parameter, R^2 values for $RVE_{MSC,ip}$ parameter predictions are shown for each input feature combination in Table 1. Each blue point represents the CNN performance during one of the 25 CNN training runs completed in a repeated five-fold CV. Each red point represents the average CNN performance from the repeated five-fold CV for a given input feature combination. The average R^2 value for the complete input set (no dropped input features) serves as a baseline estimate of CNN model performance. The baseline R^2 value is shown as a black dashed line. The percent difference between the average R^2 and the baseline is shown for each input feature combination. The results in Fig. 7 are used to compare the sensitivity of CNN predictions to different input features (discussed in Section 4.2.)

3.3. Input sampling grid study

The results of the input sampling grid study, as described in Section 2.3.2, are shown in Fig. 8. For each combination of boundary condition type and $RVE_{MSC,ip}$ parameter, R^2 values are shown for $RVE_{MSC,ip}$ parameter value predictions from CNN models trained with 3D array inputs sampled from grids extending 1, 2, 3, 4, and 5 grains ahead of the crack front. Each blue point represents the CNN performance for one of 25 training runs completed during a repeated five-fold CV. The black dashed line shows the trend in average R^2 as the distance sampled ahead of the crack front increases. The results in Fig. 8 are used to evaluate how the location of microstructural features relative to crack-front point p impact the influence of the features on $RVE_{MSC,ip}$ predictions (discussed in Section 4.3).

3.4. Saliency maps

Saliency maps (Section 2.3.3) corresponding to $d_{1,ip}$ predictions for cracks of length $a/g = 1$ under full-submodeling boundary conditions are shown in Figs. 9, 10, and 11. Due to the challenge of visualizing 3D data, select 2D slices of the 3D saliency maps are visualized. The slices are taken from the saliency map arrays at locations adjacent to a given crack-front point. The slices are oriented either parallel or perpendicular to the plane of the crack. For slices perpendicular to the crack's plane, the slices are either tangent or normal to the crack front. The saliency values within the slice are normalized to a range of 0 to 1. On each slice, the location of the crack-front point is shown as a black point and grain boundaries are shown as black lines. For reference, corresponding slices of the input E array are shown alongside the saliency maps. Although both the E and d_{cfp} arrays are used to calculate the saliency maps, only the E arrays are shown alongside the saliency maps because the E values are the larger source of variability among the two types of arrays input to the CNNs. Note, while the saliency map slices shown here comprise only a small portion of the saliency map data,[2] high-level trends observed across many saliency maps are represented.

In Fig. 9, saliency map slices from one particular microstructure (called Microstructure A) are shown. The saliency map slices are taken from five different crack-front points and are oriented parallel to the crack plane. The five crack-front points are located at 0°, 30°, 90°, 120°, and 180° along the crack front. The top row of heat maps shows the input E arrays at the slice location, while the middle row of figures shows the normalized saliency maps at the slice location. The bottom row shows the physical location of the slice relative to the crack front. On each slice, two sides of the heat maps are marked in solid and dashed green lines, respectively, to assist with visual orientation. For each slice, microstructural regions corresponding to high saliency values are labeled a, b, c, d, and e. Note that each region label is placed at the center of an area of interest within the microstructure in Microstructure A. There is overlap between sampling grids for different crack-front points, so a labeled region of microstructure may be represented in

[2] One 3D saliency map is associated with each of 14 168 data points from each of four CNN models.

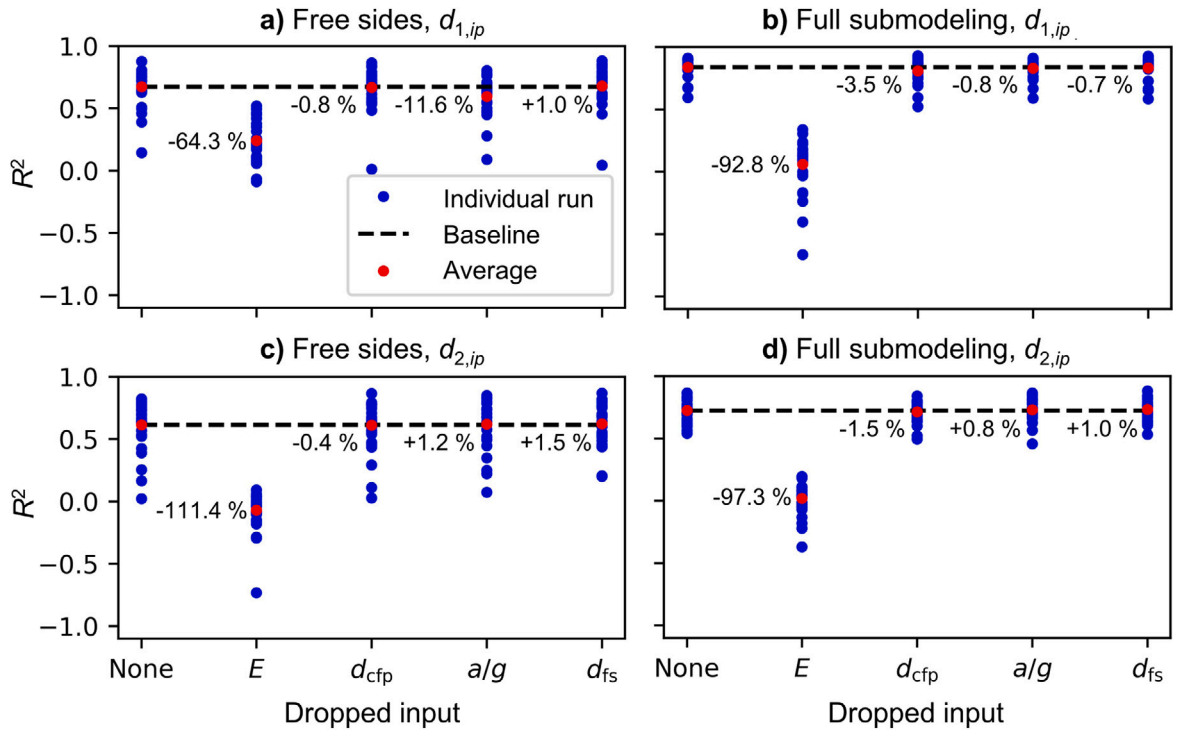

Fig. 7. Comparisons of CNN performance while predicting $RVE_{MSC,ip}$ as each input feature is removed from the CNN input set. For each input feature combination, a repeated five-fold CV is performed. The average R^2 values are compared with the baseline performance of the CNN with no features dropped from the CNN model.

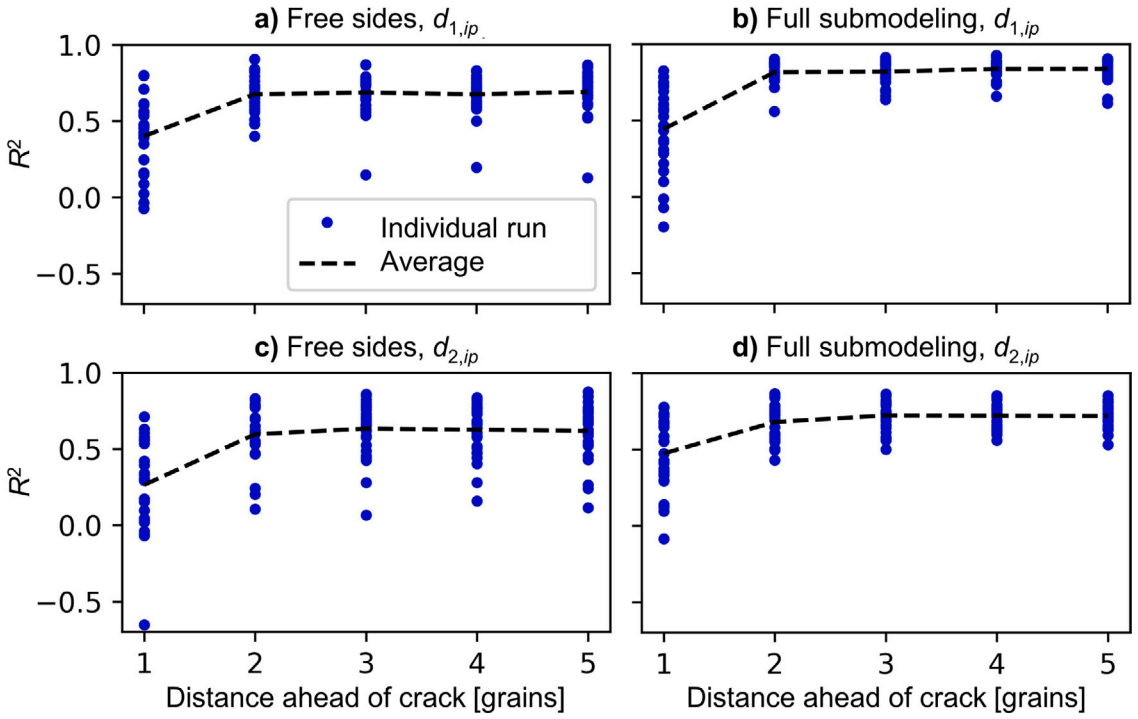

Fig. 8. Comparisons of CNN performance using 3D sampling grids of different physical dimensions. For each sampling grid, a repeated five-fold CV is performed. The R^2 values from individual training runs are shown as blue points, and average R^2 values are shown as black lines.

multiple saliency maps (i.e., regions a, b, c, d, and e). The results in Fig. 9 are used to evaluate how the sensitivity of CNN predictions to voxels corresponding to individual arrangements of grains varies at different crack-front points (discussed in Section 4.5).

Fig. 10 provides examples of saliency maps from four cracked microstructural instantiations (Microstructures B, C, D, and E). Microstructure B is a uniform microstructure case (uniform E in Fig. 5), with all grains assigned $E = 138$ GPa. Microstructure C has a random E value assigned to each grain. Microstructure D has all first-nearest-neighbor grains of a cracked grain assigned the minimum E, and Microstructure E has all second-nearest-neighbor grains of a cracked grain assigned the minimum E (2nn max/min E in Fig. 5). The saliency map slices are oriented in three ways: perpendicular to the crack plane/tangent to the crack front (Fig. 10a), parallel to the crack plane

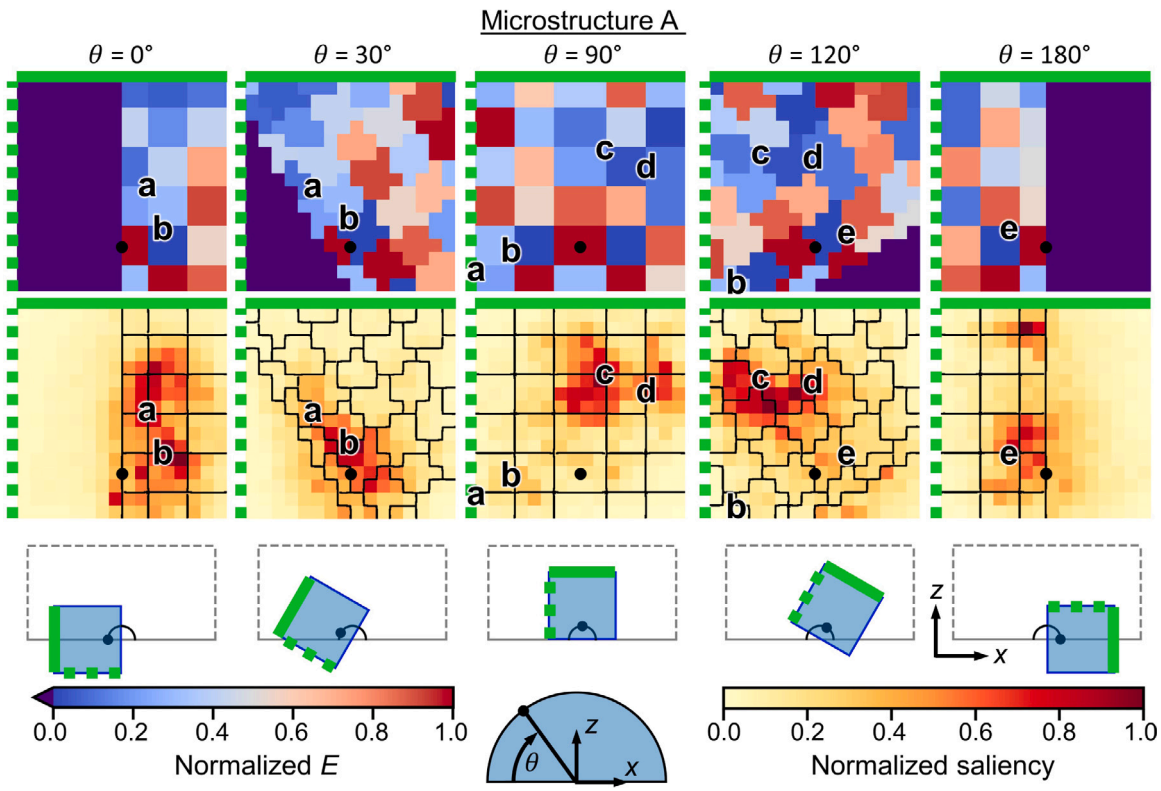

Fig. 9. Slices of E input arrays (top row) and saliency maps (middle row) from five crack-front points in Microstructure A. The location of each slice relative to the crack front is indicated in the bottom row. Regions of microstructure corresponding to voxels with high saliency values are labeled a–e; note, the regions appear in multiple slices due to overlap in the sampling grids. The black dots represent the location of the crack-front point of interest.

(Fig. 10b), and perpendicular to the crack plane/normal to the crack front (Fig. 10c). The orientations of the slices are shown in the far left column. The E input array (top row) and normalized saliency (bottom row) heat maps are shown for each slice orientation of each microstructure. As in Fig. 9, solid and dashed green lines assist with visual orientation of the heat maps relative to the crack front. The results in Fig. 10 are used to evaluate how the sensitivity of CNN predictions to individual voxels varies with variations in local microstructure (discussed in Sections 4.4 and 4.5).

Fig. 11 shows saliency maps from two additional cracked microstructural instantiations (Microstructures F and G) extracted from five independently trained CNN models. During these training runs, the data points used for training and validation are varied. Additionally, the initial CNN model weights vary from one training run to another. In the left column, a slice of the input E array, oriented parallel to the crack plane, is shown for each microstructure. The location of the slice relative to the crack is shown for reference. The five rightmost columns show saliency map slices taken at the indicated location, where each column corresponds to one of five independent CNN training runs. Certain voxel regions located near high saliency values in one or more training runs are labeled a, b, c, d, and e in the E array and saliency map slices. The results in Fig. 11 are used to evaluate how the sensitivity of CNN predictions to individual voxels varies from one trained CNN model to another (discussed in Section 4.6).

4. Discussion

4.1. Trends captured in CNN $RVE_{MSC,ip}$ predictions and corresponding $RVE_{MSC,i}$ estimates

Using geometrical (a/g, d_{fs}, d_{cfp}) and microstructural (E) inputs, four CNNs (one for each combination of boundary condition type and $RVE_{MSC,ip}$ parameter) can learn to predict $RVE_{MSC,ip}$ sizes and subsequently estimate $RVE_{MSC,i}$ sizes. Considering all crack sizes, boundary conditions, and microstructural instantiations (including the uniform microstructure case with the same E value assigned to each grain), the CNN-based estimates of $RVE_{MSC,i}$ sizes range from 8 to 3288 grains. This range of $RVE_{MSC,i}$ sizes is similar to the range of actual $RVE_{MSC,i}$ sizes (2 to 3439 grains). Fig. 4 shows how similar ranges of actual and CNN-estimated $RVE_{MSC,i}$ sizes are achieved as the CNNs learn to match key trends from the actual $RVE_{MSC,i}$ parameter values. First, the ranges of actual and CNN-estimated $d_{1,i}$ values decrease with increasing crack size, while the ranges of actual and CNN-estimated $d_{2,i}$ values remain similar among all crack sizes. Second, the actual and CNN-estimated $d_{1,i}$ values for free-sides boundary conditions tend to increase with increasing crack size. In contrast, the actual and CNN-estimated $d_{2,i}$ values for free-sides boundary conditions tend to decrease with increasing crack size. Thus, the CNNs are able to estimate reasonable ranges of $RVE_{MSC,i}$ parameters and match key trends in the actual $RVE_{MSC,i}$ data.

The CNNs can also predict $RVE_{MSC,ip}$ parameter values that match key trends in the actual $RVE_{MSC,ip}$ data. Fig. 6a shows two general trends captured by both the actual and predicted $d_{1,ip}$ values. First, the predicted $d_{1,ip}$ values lie within the same general range as the corresponding actual $d_{1,ip}$ values. Second, the actual and predicted $d_{1,ip}$ values are both generally larger at $\theta < 90°$ than at $\theta > 90°$. While not representative of all $RVE_{MSC,ip}$ predictions, Fig. 6a shows that the CNNs are able to estimate reasonable ranges of $RVE_{MSC,ip}$ parameters for a given microstructural instantiation and can capture local variations of $RVE_{MSC,ip}$ parameters along a crack front.

4.2. Relative importance of microstructural and geometrical features on volume requirement predictions

As discussed in Section 4.1, microstructural (E) and geometrical (a/g, d_{fs}, d_{cfp}) inputs to the CNN allow the CNNs to learn key trends

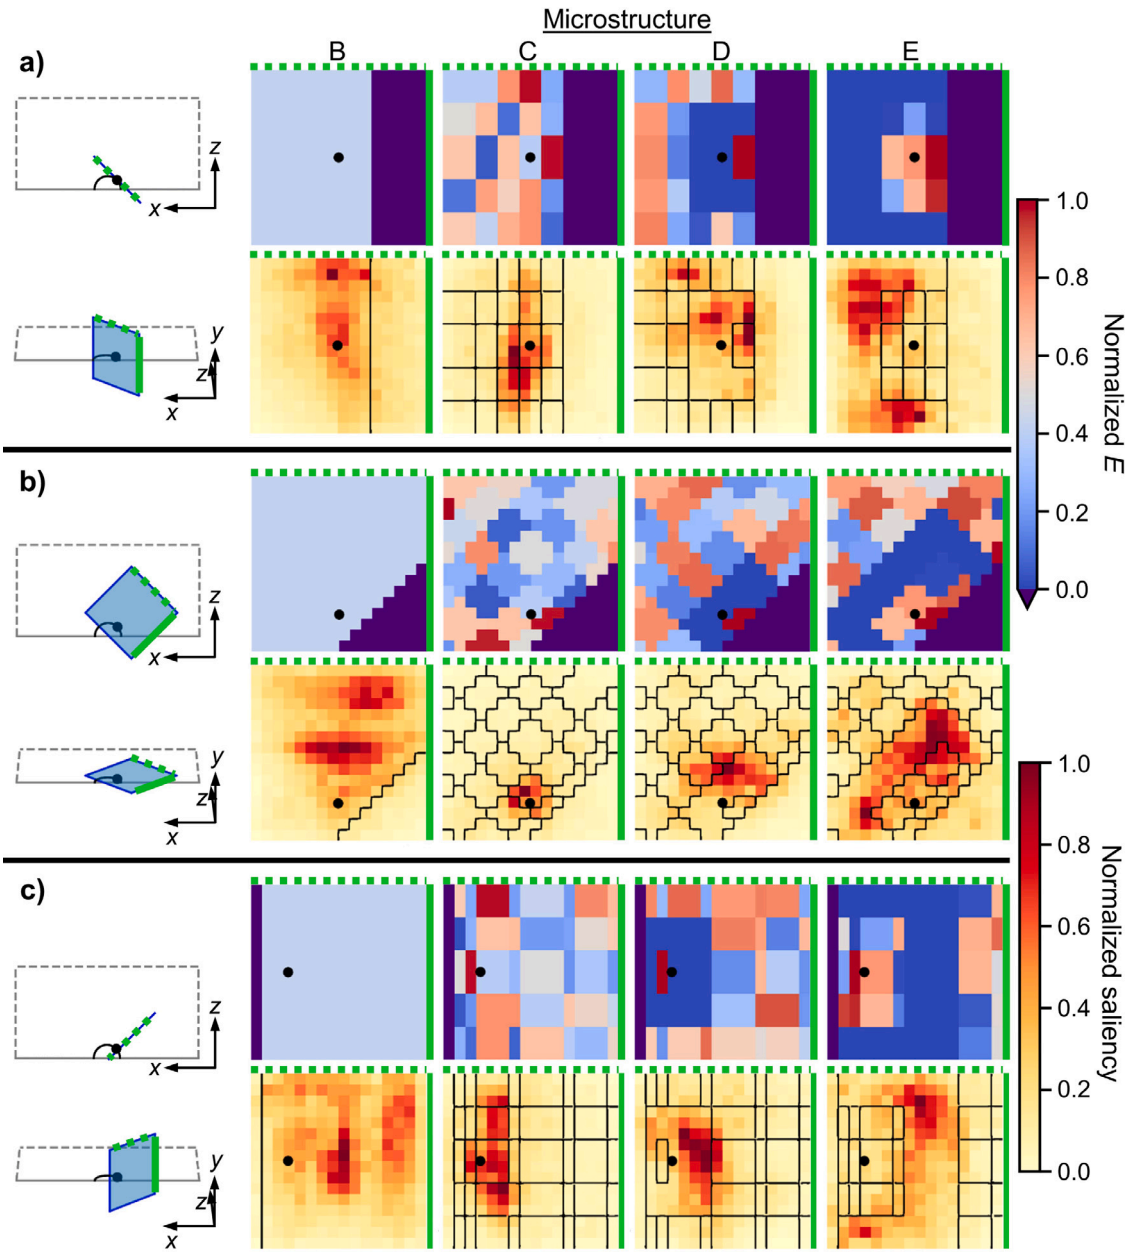

Fig. 10. Slices of E input arrays (top row) and saliency maps (bottom row) from Microstructures B, C, D, and E. The slices are oriented: (a) perpendicular to the crack plane and tangent to the crack front, (b) parallel to the crack plane, and (c) perpendicular to the crack plane and perpendicular to the crack front. The black dots represent the location of the crack-front point of interest.

in volume requirements ($\text{RVE}_{\text{MSC},i}$ and $\text{RVE}_{\text{MSC},ip}$). However, the importance of an individual input feature to the CNN predictions is not immediately evident. The sensitivity analysis, with results shown in Fig. 7, allows the relative importance of the four input features to be compared. The largest decrease in CNN performance occurs when the E input array is dropped from the CNN input set, resulting in a decrease in R^2 between 64.3 % and 111.4 % relative to the baseline predictions for $\text{RVE}_{\text{MSC},ip}$ parameters. Thus, the microstructural data captured in the E array are more critical to CNN predictions of $\text{RVE}_{\text{MSC},ip}$ than the geometrical data included in the other three inputs.

Besides the decreases in CNN performance with the removal of the E input, the largest reduction in R^2 from the baseline (11.6 %) occurs when a/g is removed from the CNN that predicts $d_{1,ip}$ for free-sides boundary conditions (Fig. 7a). The decrease in CNN performance with the removal of a/g as input is smaller for the three other combinations of boundary condition type and $\text{RVE}_{\text{MSC},ip}$ parameter (Figs. 7b,c,d). Looking at the actual and CNN-estimated $\text{RVE}_{\text{MSC},i}$ parameters in Fig. 4, the free sides, $d_{1,i}$ values (Fig. 4a) show a strong dependence on a/g, as the minimum value of $d_{1,i}$ increases with increasing crack size. For the other three cases (Figs. 4b,c,d), the minimum value of $d_{1,i}$ or $d_{2,i}$ remains less than one grain across all crack sizes. Thus, the CNNs are capturing the trends in the actual data used for training, which show that the relationship between volume requirements for MSCs and a/g varies with boundary condition type and volume parameter [15].

In the sensitivity study results (Fig. 7), the CNN inputs associated with the smallest decreases in R^2 (relative to the baseline) are the d_{cfp} array and the d_{fs} scalar. The slight impact of removing these inputs from the CNN suggests that these inputs may be unnecessary or redundant. The d_{cfp} array is likely unnecessary because the sampling grids are fixed in size when sampling inputs for a given CNN. As a result, the d_{cfp} array only varies when a sampling grid has points outside the physical volume (see Section 2.2.1). The d_{fs} input is likely

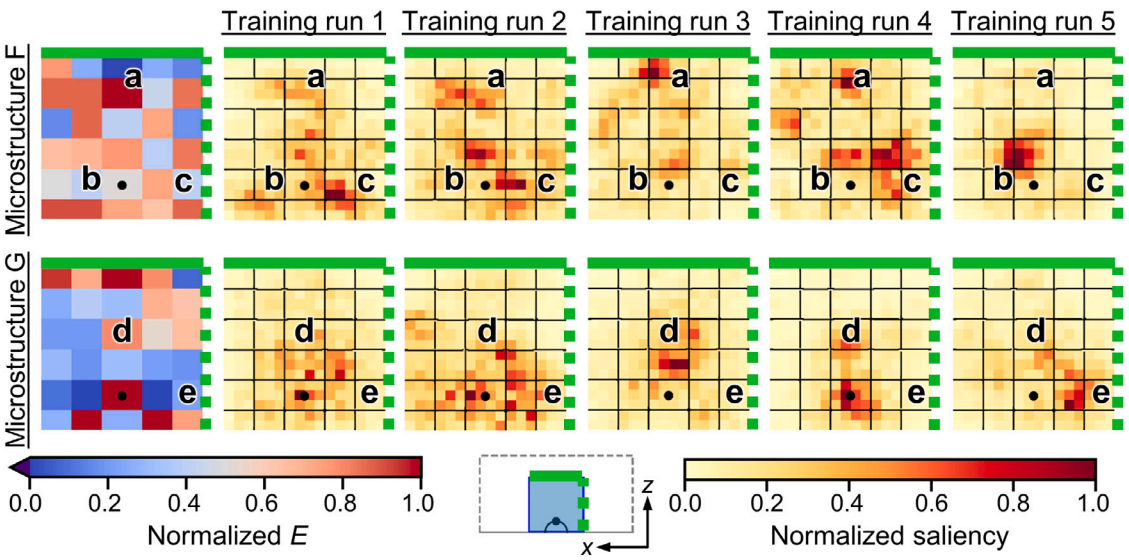

Fig. 11. Slices of E input arrays (far left column) and saliency maps (right columns) from Microstructures F and G. Each column of saliency map slices is obtained from one of five CNN training runs. The black dots represent the location of the crack-front point of interest. Locations near high saliency values are labeled a–e.

redundant because the sampling of the input arrays implicitly accounts for the location of the free surface by including values of −1 for points that fall outside the microstructural volume. Thus, the CNN may be able to identify crack-front points that are close to the free surface through the presence of placeholder (−1) values in the input arrays.

4.3. Impact of microstructural feature location on volume requirement predictions

Having established that the E array is the most critical input feature, the impact of the location of microstructural features within the E array on predictions of $\text{RVE}_{\text{MSC},ip}$ can be seen in the results of the input sampling grid study. As shown in Fig. 8, the largest increase in average R^2 occurs between input grids that extend $d_{\text{ahead}} = 1$ and $d_{\text{ahead}} = 2$ grains ahead of the crack front. Modest increases in average R^2 occur as the input grid sampling region is extended from $d_{\text{ahead}} = 2$ to $d_{\text{ahead}} = 3$ grains. Extending the input sampling grids beyond $d_{\text{ahead}} = 3$ grains provides little improvement in the average R^2. The input sampling grids with $d_{\text{ahead}} = 2$ and $d_{\text{ahead}} = 3$ grains sample 1.5 grains and 2 grains, respectively, in the direction tangent to the crack front and above/below the crack plane. Thus, most of the important microstructural features for $\text{RVE}_{\text{MSC},ip}$ predictions are within 2 grains of the crack front. However, some important microstructural features for the CNN predictions are between 2 and 3 grains away from the crack front. These findings are similar to results from our previous RVE_{MSC} study [15], where features from the second or third nearest-neighbor grains of a crack were found to be most strongly correlated with $\text{RVE}_{\text{MSC},i}$ values when features from first through fourth nearest-neighbor grains were considered. Other studies have also found that local material responses, including stress, strain, and fatigue indicator parameters, are primarily influenced by grains up to three grains away from a point of interest [55–57].

4.4. Identification of microstructural arrangements corresponding to extreme volume requirement predictions

With the goal of identifying key microstructural arrangements related to volume requirement predictions, trends in CNN-estimated $\text{RVE}_{\text{MSC},i}$ parameter values are considered. In Fig. 5, the maximum CNN-estimated $\text{RVE}_{\text{MSC},i}$ parameters for each crack size correspond to microstructures in which all second-nearest neighbors of a cracked grain are assigned the maximum or minimum E value (i.e., 2nn max/min E). These regions of extreme E values represent regions featuring elevated elastic heterogeneity. On the other hand, the minimum CNN-estimated $\text{RVE}_{\text{MSC},i}$ parameters correspond to uniform microstructures, in which all grains were assigned the same E value (i.e., uniform E). The uniform E microstructures implicitly represent a homogeneous material with no elastic heterogeneity. Although not shown, the trend of 2nn max/min E and uniform E cases having maximum and minimum CNN-estimated $\text{RVE}_{\text{MSC},i}$ parameters, respectively, is also seen for the free-sides $d_{2,i}$ and full-submodeling $d_{1,i}$ cases. Additionally, Fig. 6b shows that the microstructure surrounding the crack-front point at θ_{crit}, whose $d_{1,ip}$ prediction governs the maximum CNN-estimated $d_{1,i}$ value (for $a/g = 1$ under full-submodeling boundary conditions), contains the cluster of second-nearest neighbors assigned the minimum E. Thus, the CNN has learned to predict minimum volume requirements for microstructures with no elastic heterogeneity (uniform E) and maximum volume requirements for microstructures with clusters of large differences in local elastic properties (2nn max/min E). This is similar to the findings that increases in anisotropy (i.e., higher possible differences in local elastic properties) increased the size of an RVE for elastic homogenization of polycrystals [58–60].

The sensitivity of $\text{RVE}_{\text{MSC},ip}$ predictions to clusters of minimum E value grains is further shown in the saliency maps in Fig. 10. Microstructures D and E have regions of minimum E values at the first and second nearest-neighbor grains, respectively, of a cracked grain (which is assigned the maximum E). In the saliency map slices for Microstructure D, elevated saliency values correspond to voxels located around the first nearest-neighbor grains (dark blue grains in microstructure D). In the saliency map slices for Microstructure E, elevated saliency values are located around the second nearest-neighbor grains (dark blue grains in microstructure E). Although not shown here, similar trends within saliency maps are seen for instantiations having regions of grains with the maximum E value. Thus, at crack-front points near clusters of minimum- or maximum-valued E grains, the $\text{RVE}_{\text{MSC},ip}$ predictions are highly dependent on the grains within the minimum or maximum E clusters.

4.5. Sensitivity of volume requirement predictions to individual input voxels

Although some key microstructural arrangements have been identified using inspection of CNN-estimated $\text{RVE}_{\text{MSC},i}$ parameters, these arrangements are only present in a portion of the E arrays provided to the CNNs. Saliency maps are visually inspected to compare the sensitivity of CNN predictions to voxels corresponding to other microstructural

arrangements in the E arrays. Visual (qualitative) inspection represents a common state of practice for the evaluation of CNN saliency maps in machine learning, and further interpretation of saliency maps using quantitative measures remains an open area of research, even within the computer science community [61]. Recall that saliency maps compare the relative importance of individual voxels to CNN output within the input arrays. For this reason, the saliency map slices have been individually normalized such that saliency values in a given slice vary from 0 to 1. While this normalization technique does not allow for the comparison of saliency values among different slices, it does allow the influence of individual voxels on CNN output to be easily compared and visualized within a slice. As discussed in Section 3.4, the saliency maps presented in Figs. 9 and 10 do not represent all saliency data for the CNNs, as a 3D saliency map exists for every crack-front point. However, the included saliency map slices capture general trends observed within the saliency data. Several key findings from the saliency maps are discussed in the following paragraphs.

First, the saliency maps show that CNN predictions are more influenced by the values of voxels located within the physical volume of microstructure than voxels located outside of the physical volume. In Figs. 9 and 10, voxels located outside the physical volume are shown in purple in the E array slices. The saliency values at voxels outside the physical volume are low relative to other saliency values in the slice. In some cases, sharp decreases in saliency values are observed when crossing from inside the physical volume to outside the physical volume. A good example of this is seen in the saliency map slice in Fig. 9 for $\theta = 0°$, where strong saliency values are seen just inside the physical volume, and much lower saliency values are seen outside of the physical volume. Thus, the saliency maps demonstrate that the CNN has learned to focus on data corresponding to regions inside the physical volume.

Second, the saliency maps show that the locations of voxels (relative to a crack-front point of interest) that most influence $\text{RVE}_{\text{MSC},ip}$ predictions vary at different crack-front points and in different microstructural instantiations. The saliency-map slices shown in Fig. 9 exemplify how the locations of voxels having the most influence on the CNN's learning trends for $\text{RVE}_{\text{MSC},ip}$ vary at different crack-front points within the same microstructural configuration. For example, voxels to the left of the crack-front point have high saliency values at $\theta = 0°$, but have low saliency values at $\theta = 180°$. The saliency maps shown in Fig. 10 demonstrate how the influential microstructural regions around a fixed crack-front position vary with respect to microstructural instantiation. For each of the three slice orientations, the location of voxels having high saliency values varies in each of the four microstructures. Thus, rather than looking at fixed voxel locations within the E array, the CNN has learned to identify arrangements of grains that tend to influence predictions of J-integral convergence behavior (with respect to volume size).

Third, the saliency maps show that multiple $\text{RVE}_{\text{MSC},ip}$ predictions can be influenced by one arrangement of grains. Examples of this are shown in Fig. 9. Recall, the labels in Fig. 9 correspond to regions of elevated saliency values within Microstructure A and are represented in saliency maps for multiple crack-front points, as there is overlap between sampling grids for different crack-front points. In the saliency maps for crack-front points at $\theta = 0°$ and $\theta = 30°$, regions a and b feature high saliency values, despite a shift in the location of the voxels corresponding to regions a and b within the two arrays. A similar trend is seen with the saliency values for regions c and d in the saliency maps at $\theta = 90°$ and $\theta = 120°$. Thus, the CNN has learned that predictions of volume requirements for multiple crack-front points may be influenced by the same arrangement of grains.

Fourth, the saliency maps show that the influence of a given arrangement of grains on $\text{RVE}_{\text{MSC},ip}$ predictions varies along a crack front. In Fig. 9, region b corresponds to high saliency values at $\theta = 0°$ and $\theta = 30°$. The normalized saliency values around region b are only slightly elevated at $\theta = 90°$, then decrease to nearly zero at $\theta = 120°$. Similarly, the normalized saliency values around region a are high at $\theta = 0°$, lower at $\theta = 30°$, then nearly zero at $\theta = 90°$; and the normalized saliency values around region e are slightly elevated at $\theta = 120°$ and high at $\theta = 180°$. Thus, the CNN has learned that the influence of a certain arrangement of grains on the prediction of volume requirements varies from one crack-front point to another.

4.6. Considerations for the interpretation of saliency maps

The interpretation of the CNNs via saliency maps poses several challenges. First, the saliency maps comprise a large set of 3D data. As described in Section 2.3.3, a 3D saliency map is associated with each crack-front point for each of the four CNN models. In total, the saliency map data contain four sets of 14 168 $16 \times 16 \times 16$ arrays. Second, saliency maps can vary from one training run to another. Fig. 11 shows variations in saliency map slices from Microstructures F and G during five different CNN training runs. Third, saliency maps are primarily visual interpretation tools that have inherent limitations and cannot be easily interpreted quantitatively [61–64]. Thus, users must manually inspect many saliency maps to identify general, model-level trends related to CNN learning. Additionally, a user may need domain knowledge to identify and interpret key trends in the saliency maps.

This work takes steps to overcome the challenges associated with saliency maps. First, only slices adjacent to the crack-front point of interest are manually inspected, as 3D saliency maps are difficult to visualize. Second, tens of thousands of saliency maps are available, so only saliency map slices for $d_{1,ip}$ predictions under full-submodeling boundary conditions with crack size $a/g = 1$ are considered when initially identifying trends in the saliency maps. Domain knowledge related to the $\text{RVE}_{\text{MSC},ip}$ data set is leveraged in identifying trends. Once trends are identified, other sets of saliency maps are inspected to verify that the identified trends exist in other sets of saliency maps.

Saliency maps are analyzed qualitatively to avoid issues with variations in saliency maps during different training runs. In the qualitative analysis of the saliency maps, the identification of high-level trends is prioritized over the inspection of individual saliency values. As seen in Fig. 11, the location of maximum saliency varies between different training runs. For example, In Microstructure F, the maximum saliency value is located near a, b, or c depending on the training run. Similarly, in Microstructure G, the maximum saliency value is around the crack-front point, d, or e depending on the training run. Due to the local fluctuations, emphasis is placed on identifying general regions of elevated saliency values during the interpretation of the saliency maps. In Microstructure F, elevated saliency values tend to be located near points a, b, and c in all training runs. Training runs 3 and 4 also have an additional region of elevated saliency values on the left edge of the saliency maps for Microstructure F. In Microstructure G, elevated saliency values tend to be located near the crack-front point, d, or e. Thus, qualitative analysis of elevated saliency regions shows that the CNNs learn similar patterns during different training runs, despite local fluctuations between training runs.

4.7. Implications on the determination of RVE_{MSC}

As discussed in Section 4.1, the CNNs capture general trends in $\text{RVE}_{\text{MSC},i}$ and $\text{RVE}_{\text{MSC},ip}$, but some observations in the actual $\text{RVE}_{\text{MSC},i}$ and $\text{RVE}_{\text{MSC},ip}$ parameter values are not captured in the CNN predictions or estimates. First, there are noticeable differences between actual and CNN-estimated $\text{RVE}_{\text{MSC},i}$ parameter ranges. For example, the ranges of CNN-estimated $d_{1,i}$ values in Fig. 4a do not cover the full range of actual $d_{1,i}$ values for $a/g = 0.25$, 0.45, and 1.0. The maximum CNN-estimated $d_{2,i}$ values for $a/g = 1.0$ in Figs. 4c,d exceed the maximum actual $d_{2,i}$ values by more than two grains. Second, the predicted $d_{1,ip}$ values in Fig. 6a deviate from some trends that are present in the actual $d_{1,ip}$ values. The actual $d_{1,ip}$ values for the maximum $\text{RVE}_{\text{MSC},i}$ instantiation increase between crack-front points located at $\theta = 150°$

and 180°, while the predicted $d_{1,ip}$ values do not increase. These differences between actual and predicted (or CNN-estimated) $RVE_{MSC,i}$ (or $RVE_{MSC,ip}$) parameters support the finding from Part I [16] that CNN predictions are not sufficiently accurate to completely replace FE simulations when determining RVE_{MSC}.

The findings discussed in the previous sections also demonstrate that local microstructural arrangements strongly contribute to variations in both actual and CNN-estimated $RVE_{MSC,i}$ values. The demonstrated dependence of $RVE_{MSC,i}$ on local microstructure emphasizes the importance of considering a variety of microstructural instantiations in the determination of RVE_{MSC}. Overall, the goal of determining RVE_{MSC} is to establish conservative guidelines for volume requirements in studies of MSCs. The conservative guideline should guarantee convergence of crack-front parameters (J-integral values in this work) with respect to volume size in arbitrary microstructural instantiations. In other words, RVE_{MSC} should be equivalent to the maximum $RVE_{MSC,i}$ size identified among a population of microstructures. Given the strong dependence of $RVE_{MSC,i}$ on local microstructure, determining a conservative RVE_{MSC} size requires the identification of $RVE_{MSC,i}$ sizes for a variety of microstructural instantiations.

4.8. Implications for applications of CNN interpretation to computational materials science problems

Recently, CNNs have been implemented in computational materials science research because of their ability to make rapid material behavior predictions based on high-dimensional material data (viz., microstructural images). However, CNN models are "black-box" ML models. The methods explored in this work demonstrate how insight into "black-box" CNN learning can be gained through interpretation. An inspection of CNN predictions reveals trends learned by a CNN. For example, an input feature sensitivity analysis (Section 2.3.1) reveals the most critical input features to material behavior predictions. An input sampling grid study (Section 2.3.2) reveals the region of local material features that primarily influences predictions of a target material behavior. Saliency maps (Section 2.3.3) compare the relative influence of different input voxels in predicting material behavior.

It is important to note that findings from CNN interpretation represent trends learned by the CNN, but not necessarily physics-based trends. Consider a comparison of the trends in actual and CNN-estimated $RVE_{MSC,i}$ parameter values. In Fig. 5b, the uniform E microstructures have the minimum actual and CNN-estimated $d_{2,i}$ values. In this case, a trend learned by the CNN matches a trend in the actual data. In Fig. 5a, the uniform E microstructures have the minimum CNN-estimated $d_{1,i}$ values. However, the uniform E microstructures do not always have the minimum actual $d_{1,i}$ values. Specifically, at crack sizes of $a/g = 1.0$ and $a/g = 3.0$, the minimum $d_{1,i}$ values for free-sides boundary conditions are associated with microstructures other than the uniform E microstructure. The trend of predicting minimum volume requirements for uniform E cases does not reflect the trend in the actual data. Thus, the trends learned by a CNN cannot be guaranteed to represent the physics of a problem. Nonetheless, trends learned by a CNN could be used to inform the design of future experiments investigating the physics of the problem.

5. Limitations

The conclusions drawn from the interpretation of CNNs are limited to the assumptions made when obtaining the training data [15]. The training data were limited to results from previous linear-elastic FE simulations of idealized microstructural instantiations, where each instantiation comprised cube-shaped grains and contained one semi-circular surface crack. Grain orientations were implicitly represented by varying the E assigned to each grain, and the crack was assumed to be static. Future studies will be needed to incorporate more complex conditions into the determination of RVE_{MSC}. However, the findings from Parts I [16] and II of this work can be applied to improve the tractability and interpretability of future RVE_{MSC} studies through the assistance of CNN predictions.

6. Conclusions

Convolutional neural networks (CNNs) are a type of machine learning (ML) model capable of handling high-dimensional data and making highly accurate predictions when provided with sufficient training data. While CNNs have grown in popularity among the materials science community, they are considered "black-box" models in that the relationships between inputs and model predictions are not directly interpretable. In this work, a variety of interpretation methods are implemented to uncover the learning trends of CNN models trained on the task of predicting minimum volume requirements for microstructures containing a microstructurally small crack (MSC).

In Part I [16], CNN models were trained to estimate RVE_{MSC}, or the size of a representative volume element for a microstructure containing an MSC, which, by definition, is the minimum microstructural volume required to achieve convergence of crack-front parameters with respect to volume size [15]. Training data were derived from computationally expensive finite-element (FE) simulations. CNNs were implemented to expedite the process of determining RVE_{MSC}, specifically by predicting $RVE_{MSC,ip}$ (the minimum volume of microstructure required around an MSC for the J-integral value at crack-front point p in microstructural instantiation i to converge with respect to volume size) given microstructural and geometrical input features. The predictions of $RVE_{MSC,ip}$ parameters ($d_{1,ip}$ and $d_{2,ip}$) were subsequently used to make CNN-based estimates of $RVE_{MSC,i}$ (the minimum volume of microstructure required around an MSC for the J-integral values in microstructural instantiation i to converge with respect to volume size) for a given microstructural instantiation and RVE_{MSC} for an entire population of microstructural instantiations.

In this work (Part II), several approaches are used to interpret the CNNs discussed in Part I. First, CNN predictions are inspected to identify high-level relationships between microstructure and predicted volume requirements. Second, an input feature sensitivity study compares the importance of different input features to CNN predictions. Third, the physical dimensions of the input sampling grids are varied to determine how many nearest-neighbor grains influence CNN predictions of volume requirements. Fourth, saliency maps are inspected to explore the influence of individual microstructural arrangements on volume requirement predictions. The following conclusions are drawn:

- Inspection of CNN predictions (or estimates) reveals certain arrangements of grains that consistently correspond to extreme (maximum or minimum) predictions of volume requirements. There are some cases for which grain arrangements predicted by the CNNs to have extreme-volume requirements do not actually have extreme-volume requirements. However, in most cases, the CNN-predicted extreme-volume arrangements agree with those identified in the training data. Thus, trends in the CNN predictions often, but not always, correspond to physics-based trends.
- Given microstructural and geometrical inputs, CNNs can capture many trends in $RVE_{MSC,i}$, including variations due to crack size and microstructure. Through an input feature sensitivity study, microstructural input is shown to be more important to CNN predictions of $RVE_{MSC,ip}$ than geometrical input. The importance of crack size to volume requirements depends on the $RVE_{MSC,ip}$ parameter and boundary conditions for which $RVE_{MSC,ip}$ is being predicted.
- Varying the physical size of the 3D sampling grids for the CNN input arrays shows that most of the important microstructural features to $RVE_{MSC,ip}$ predictions are within two to three grains of a crack-front point.
- Saliency maps extracted from trained CNNs demonstrate the ability of a CNN to identify patterns of specific microstructural arrangements within a given volume that tend to influence predictions of the convergence behavior of crack-front parameters. For example, voxels corresponding to second-nearest-neighbor grains

(of a cracked grain) that have minimum or maximum E values consistently have high saliency values. On the other hand, voxels sampled from outside the physical volume of microstructure consistently have low saliency values, demonstrating the CNN's ability to de-emphasize irrelevant regions.

- The strong dependence of $RVE_{MSC,ip}$ predictions (and thus, CNN-based $RVE_{MSC,i}$ estimates) on microstructure highlights the need to define a conservative RVE_{MSC} size by considering multiple microstructural instantiations. By considering many microstructural instantiations and identifying the upper bound of volume requirements, the resulting RVE_{MSC} guidelines can generalize to a wide range of microstructural instantiations.

The interpretation approaches explored in this work demonstrate that "black-box" CNN models can provide benefits beyond the time-savings capabilities demonstrated in Part I [16]. The results show that analyses of trained CNNs can provide model-verification approaches for determining whether a model is learning expected trends. The results also show that CNN interpretation can be used to compare input feature importance and identify key patterns within inherently complex data. As demonstrated in Part I [16] and in this work, CNNs can be harnessed to both expedite the determination of and provide potential insights into RVE_{MSC}.

Given the complex, high-dimensional nature of materials data, the interpretation of CNNs is a potentially valuable tool for better understanding material behavior in other computational materials science applications. During training, CNNs can learn to process complex, high-dimensional material data to extract relevant, low-dimensional features related to material behavior trends. Interpretation methods can then be used to explore the features (or patterns) identified by the CNN to be most relevant to the predictions. Future studies in computational materials science can harness the data processing power of CNNs alongside CNN interpretation to explore trends in complex, high-dimensional materials data.

CRediT authorship contribution statement

Karen J. DeMille: Conceptualization, Methodology, Software, Formal analysis, Data curation, Writing – original draft, Visualization, Funding acquisition. **Ashley D. Spear:** Conceptualization, Resources, Writing – review & editing, Supervision, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The processed data required to reproduce these findings are available to download from Mendeley Data at https://doi.org/10.17632/nrntzs2zb7.1.

Acknowledgments

KJD was supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate (NDSEG) Fellowship Program. This material is based upon work supported by the National Science Foundation under Grant No. CMMI-1752400. The support and resources from the Center for High Performance Computing at the University of Utah are gratefully acknowledged. The authors gratefully acknowledge Aowabin Rahman and Vignesh Babu Rao for their valuable discussions and support.

References

[1] S. Suresh, Fatigue of Materials, Cambridge University Press, 1998, http://dx.doi.org/10.1017/CBO9780511806575.
[2] D.L. McDowell, Damage mechanics and metal fatigue: A discriminating perspective, Int. J. Damage Mech. 8 (4) (1999) 376–403, http://dx.doi.org/10.1177/105678959900800406.
[3] W. Ludwig, J.-Y. Buffière, S. Savelli, P. Cloetens, Study of the interaction of a short fatigue crack with grain boundaries in a cast Al alloy using X-ray microtomography, Acta Mater. 51 (3) (2003) 585–598, http://dx.doi.org/10.1016/S1359-6454(02)00320-8.
[4] A. Ayyar, N. Chawla, Microstructure-based modeling of the influence of particle spatial distribution and fracture on crack growth in particle-reinforced composites, Acta Mater. 55 (18) (2007) 6064–6073, http://dx.doi.org/10.1016/j.actamat.2007.06.044.
[5] A.R. Najafi, A.R. Arshi, M.R. Eslami, S. Fariborz, M.H. Moeinzadeh, Micromechanics fracture in osteonal cortical bone: A study of the interactions between microcrack propagation, microstructure and the material properties, J. Biomech. 40 (12) (2007) 2788–2795, http://dx.doi.org/10.1016/j.jbiomech.2007.01.017.
[6] G.M. Castelluccio, W.D. Musinski, D.L. McDowell, Recent developments in assessing microstructure-sensitive early stage fatigue of polycrystals, Curr. Opin. Solid State Mater. Sci. 18 (4) (2014) 180–187, http://dx.doi.org/10.1016/j.cossms.2014.03.001.
[7] Y. Takahashi, T. Shikama, S. Yoshihara, T. Aiura, H. Noguchi, Study on dominant mechanism of high-cycle fatigue life in 6061-T6 aluminum alloy through microanalyses of microstructurally small cracks, Acta Mater. 60 (6–7) (2012) 2554–2567, http://dx.doi.org/10.1016/j.actamat.2012.01.023.
[8] W. Schaef, M. Marx, A numerical description of short fatigue cracks interacting with grain boundaries, Acta Mater. 60 (5) (2012) 2425–2436, http://dx.doi.org/10.1016/j.actamat.2012.01.013.
[9] M. Rocha, E. Brühwiler, A. Nussbaumer, Microstructural influence on the scatter in the fatigue life of steel reinforcement bars, Int. J. Fatigue 75 (2015) 205–212, http://dx.doi.org/10.1016/j.ijfatigue.2015.01.016.
[10] K. Tokaji, T. Ogawa, The growth behaviour of microstructurally small fatigue cracks in metals, in: K.J. Miller, E.R. de Los Rios (Eds.), Short Fatigue Cracks, ESIS 13, Mechanical Engineering Publications, London, 1992, pp. 85–99.
[11] S. Suresh, R.O. Ritchie, Propagation of short fatigue cracks, Int. Met. Rev. 29 (1) (1984) 445–475, http://dx.doi.org/10.1179/imtr.1984.29.1.445.
[12] K.S. Chan, J. Lankford, The role of microstructural dissimilitude in fatigue and fracture of small cracks, Acta Metall. 36 (1) (1988) 193–206, http://dx.doi.org/10.1016/0001-6160(88)90038-7.
[13] K. Tokaji, T. Ogawa, Y. Harada, The growth of small fatigue cracks in a low carbon steel; The effect of microstructure and limitations of linear elastic fracture mechanics, Fatigue Fract. Eng. Mater. Struct. 9 (3) (1986) 205–217, http://dx.doi.org/10.1111/j.1460-2695.1986.tb00447.x.
[14] R.O. Ritchie, J. Lankford, Small fatigue cracks: A statement of the problem and potential solutions, Mater. Sci. Eng. 84 (1986) 11–16, http://dx.doi.org/10.1016/0025-5416(86)90217-X.
[15] K.J. DeMille, A.D. Spear, Determination of representative volume elements for small cracks in heterogeneous, linear-elastic domains, Eng. Fract. Mech. 220 (2019) 106643, http://dx.doi.org/10.1016/j.engfracmech.2019.106643.
[16] K.J. DeMille, A.D. Spear, Convolutional neural networks for expediting the determination of minimum volume requirements for studies of microstructurally small cracks, Part I: Model implementation and predictions, Comput. Mater. Sci. 207 (2022) 111290, http://dx.doi.org/10.1016/j.commatsci.2022.111290.
[17] A. Cecen, H. Dai, Y.C. Yabansu, S.R. Kalidindi, L. Song, Material structure-property linkages using three-dimensional convolutional neural networks, Acta Mater. 146 (2018) 76–84, http://dx.doi.org/10.1016/j.actamat.2017.11.053.
[18] A. Beniwal, R. Dadhich, A. Alankar, Deep learning based predictive modeling for structure-property linkages, Materialia 8 (2019) 100435, http://dx.doi.org/10.1016/j.mtla.2019.100435.
[19] Z. Yang, Y.C. Yabansu, R. Al-Bahrani, W.-k. Liao, A.N. Choudhary, S.R. Kalidindi, A. Agrawal, Deep learning approaches for mining structure-property linkages in high contrast composites from simulation datasets, Comput. Mater. Sci. 151 (2018) 278–287, http://dx.doi.org/10.1016/j.commatsci.2018.05.014.
[20] X. Li, Z. Liu, S. Cui, C. Luo, C. Li, Z. Zhuang, Predicting the effective mechanical property of heterogeneous materials by image based modeling and deep learning, Comput. Methods Appl. Mech. Engrg. 347 (2019) 735–753, http://dx.doi.org/10.1016/j.cma.2019.01.005.
[21] D.W. Abueidda, M. Almasri, R. Ammourah, U. Ravaioli, I.M. Jasiuk, N.A. Sobh, Prediction and optimization of mechanical properties of composites using convolutional neural networks, Compos. Struct. 227 (2019) 111264, http://dx.doi.org/10.1016/j.compstruct.2019.111264.
[22] S. Ye, B. Li, Q. Li, H.-P. Zhao, X.-Q. Feng, Deep neural network method for predicting the mechanical properties of composites, Appl. Phys. Lett. 115 (16) (2019) 161901, http://dx.doi.org/10.1063/1.5124529.
[23] C. Herriott, A.D. Spear, Predicting microstructure-dependent mechanical properties in additively manufactured metals with machine-and deep-learning methods, Comput. Mater. Sci. 175 (2020) 109599, http://dx.doi.org/10.1016/j.commatsci.2020.109599.

[24] S.M. Azimi, D. Britz, M. Engstler, M. Fritz, F. Mücklich, Advanced steel microstructural classification by deep learning methods, Sci. Rep. 8 (2018) 2128, http://dx.doi.org/10.1038/s41598-018-20037-5.
[25] B.L. DeCost, B. Lei, T. Francis, E.A. Holm, High throughput quantitative metallography for complex microstructures using deep learning: A case study in ultrahigh carbon steel, Microsc. Microanal. 25 (1) (2019) 21–29, http://dx.doi.org/10.1017/S1431927618015635.
[26] E.A. Holm, R. Cohn, N. Gao, A.R. Kitahara, T.P. Matson, B. Lei, S.R. Yarasi, Overview: Computer vision and machine learning for microstructural characterization and analysis, Metall. Mater. Trans. A 51 (2020) 5985–5999, http://dx.doi.org/10.1007/s11661-020-06008-4.
[27] H. Kim, J. Inoue, T. Kasuya, Unsupervised microstructure segmentation by mimicking metallurgists' approach to pattern recognition, Sci. Rep. 10 (2020) 17835, http://dx.doi.org/10.1038/s41598-020-74935-8.
[28] S. Evsevleev, S. Paciornik, G. Bruno, Advanced deep learning-based 3D microstructural characterization of multiphase metal matrix composites, Adv. Eng. Mater. 22 (4) (2020) 1901197, http://dx.doi.org/10.1002/adem.201901197.
[29] A. Sciazko, Y. Komatsu, T. Shimura, N. Shikazono, Segmentation of solid oxide cell electrodes by patch convolutional neural network, J. Electrochem. Soc. 168 (4) (2021) 044504, http://dx.doi.org/10.1149/1945-7111/abef84.
[30] M. Warmuzek, M. Żelawski, T. Jałocha, Application of the convolutional neural network for recognition of the metal alloys microstructure constituents based on their morphological characteristics, Comput. Mater. Sci. 199 (2021) 110722, http://dx.doi.org/10.1016/j.commatsci.2021.110722.
[31] H. Yu, D. Guo, Z. Yan, L. Fu, J. Simmons, C.P. Przybyla, S. Wang, Weakly supervised easy-to-hard learning for object detection in image sequences, Neurocomputing 398 (2020) 71–82, http://dx.doi.org/10.1016/j.neucom.2020.02.075.
[32] Z. Yang, Y.C. Yabansu, D. Jha, W.-k. Liao, A.N. Choudhary, S.R. Kalidindi, A. Agrawal, Establishing structure-property localization linkages for elastic deformation of three-dimensional high contrast composites using deep learning approaches, Acta Mater. 166 (2019) 335–345, http://dx.doi.org/10.1016/j.actamat.2018.12.045.
[33] A. Frankel, K. Tachida, R. Jones, Prediction of the evolution of the stress field of polycrystals undergoing elastic-plastic deformation with a hybrid neural network model, Mach. Learn.: Sci. Technol. 1 (3) (2020) 035005, http://dx.doi.org/10.1088/2632-2153/ab9299.
[34] B.P. Croom, M. Berkson, R.K. Mueller, M. Presley, S. Storck, Deep learning prediction of stress fields in additively manufactured metals with intricate defect networks, Mech. Mater. 165 (2022) 104191, http://dx.doi.org/10.1016/j.mechmat.2021.104191.
[35] A. Bhaduri, A. Gupta, L. Graham-Brady, Stress field prediction in fiber-reinforced composite materials using a deep learning approach, Composites B 238 (2022) 109879, http://dx.doi.org/10.1016/j.compositesb.2022.109879.
[36] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, D. Pedreschi, A survey of methods for explaining black box models, ACM Comput. Surv. 51 (5) (2018) 93, http://dx.doi.org/10.1145/3236009.
[37] B. Wang, R. Ma, J. Kuang, Y. Zhang, How decisions are made in brains: Unpack "black box" of CNN with Ms. Pac-Man video game, IEEE Access 8 (2020) 142446–142458, http://dx.doi.org/10.1109/ACCESS.2020.3013645.
[38] M.J. Eslamibidgoli, F.P. Tipp, J. Jitsev, J. Jankovic, M.H. Eikerling, K. Malek, Convolutional neural networks for high throughput screening of catalyst layer inks for polymer electrolyte fuel cells, RSC Adv. 11 (2021) 32126–32134, http://dx.doi.org/10.1039/D1RA05324H.
[39] B.S.S. Pokuri, S. Ghosal, A. Kokate, S. Sarkar, B. Ganapathysubramanian, Interpretable deep learning for guided microstructure-property explorations in photovoltaics, Npj Comput. Mater. 5 (2019) 95, http://dx.doi.org/10.1038/s41524-019-0231-y.
[40] C. Kantzos, J. Lao, A. Rollett, Design of an interpretable Convolutional Neural Network for stress concentration prediction in rough surfaces, Mater. Charact. 158 (2019) 109961, http://dx.doi.org/10.1016/j.matchar.2019.109961.
[41] A. Rahman, P. Deshpande, M.S. Radue, G.M. Odegard, S. Gowtham, S. Ghosh, A.D. Spear, A machine learning framework for predicting the shear strength of carbon nanotube-polymer interfaces based on molecular dynamics simulation data, Compos. Sci. Technol. 207 (2021) 108627, http://dx.doi.org/10.1016/j.compscitech.2020.108627.
[42] A.H. Sung, Ranking importance of input parameters of neural networks, Expert Syst. Appl. 15 (3–4) (1998) 405–411, http://dx.doi.org/10.1016/S0957-4174(98)00041-4.
[43] K. DeMille, A. Spear, J-integral data for determining minimum volume requirements for studies of microstructurally small cracks in heterogeneous, linear-elastic domains, in: Mendeley Data, V1, 2022, http://dx.doi.org/10.17632/nrntzs2zb7.1.
[44] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, Nature 521 (2015) 436–444, http://dx.doi.org/10.1038/nature14539.
[45] K. O'Shea, R. Nash, An introduction to convolutional neural networks, 2015, arXiv preprint arXiv:1511.08458. http://dx.doi.org/10.48550/arXiv.1511.08458.
[46] V. Bennett, D.L. McDowell, Polycrystal orientation effects on microslip and mixed-mode behavior of microstructurally small cracks, in: K.J. Miller, D.L. McDowell (Eds.), Mixed-Mode Crack Behavior, ASTM International, West Conshohocken, PA, USA, 1999, pp. 203–228, http://dx.doi.org/10.1520/STP14251S.
[47] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, 2015, arXiv preprint arXiv:1409.1556. http://dx.doi.org/10.48550/arXiv.1409.1556.
[48] F. Chollet, et al., Keras, 2015, https://keras.io.
[49] M. Abadi, et al., TensorFlow: Large-scale machine learning on heterogeneous systems, 2015, URL https://www.tensorflow.org/. Software available from tensorflow.org.
[50] A. Rovinelli, M.D. Sangid, H. Proudhon, Y. Guilhem, R.A. Lebensohn, W. Ludwig, Predicting the 3D fatigue crack growth rate of small cracks using multimodal data via Bayesian networks: In-situ experiments and crystal plasticity simulations, J. Mech. Phys. Solids 115 (2018) 208–229, http://dx.doi.org/10.1016/j.jmps.2018.03.007.
[51] A. Rovinelli, M.D. Sangid, H. Proudhon, W. Ludwig, Using machine learning and a data-driven approach to identify the small fatigue crack driving force in polycrystalline materials, Npj Comput. Mater. 4 (2018) 35, http://dx.doi.org/10.1038/s41524-018-0094-7.
[52] K. Simonyan, A. Vedaldi, A. Zisserman, Deep inside convolutional networks: Visualising image classification models and saliency maps, 2014, arXiv preprint arXiv:1312.6034. http://dx.doi.org/10.48550/arXiv.1312.6034.
[53] Y. Kubota, Contributors, tf-keras-vis, 2023, GitHub. https://keisen.github.io/tf-keras-vis-docs/.
[54] D. Smilkov, N. Thorat, B. Kim, F. Viégas, M. Wattenberg, Smoothgrad: removing noise by adding noise, 2017, arXiv preprint arXiv:1706.03825. http://dx.doi.org/10.48550/arXiv.1706.03825.
[55] K.S. Stopka, M. Yaghoobi, J.E. Allison, D.L. McDowell, Simulated effects of sample size and grain neighborhood on the modeling of extreme value fatigue response, Acta Mater. (2022) 117524, http://dx.doi.org/10.1016/j.actamat.2021.117524.
[56] T. Ozturk, C. Stein, R. Pokharel, C. Hefferan, H. Tucker, S. Jha, R. John, R.A. Lebensohn, P. Kenesei, R.M. Suter, A.D. Rollett, Simulation domain size requirements for elastic response of 3D polycrystalline materials, Modelling Simul. Mater. Sci. Eng. 24 (2016) 015006, http://dx.doi.org/10.1088/0965-0393/24/1/015006.
[57] G.M. Castelluccio, D.L. McDowell, Microstructure and mesh sensitivities of mesoscale surrogate driving force measures for transgranular fatigue cracks in polycrystals, Mater. Sci. Eng. A 639 (2015) 626–639, http://dx.doi.org/10.1016/j.msea.2015.05.048.
[58] M. Nygårds, Number of grains necessary to homogenize elastic materials with cubic symmetry, Mech. Mater. 35 (11) (2003) 1049–1057, http://dx.doi.org/10.1016/S0167-6636(02)00325-3.
[59] Z.-Y. Ren, Q.-S. Zheng, Effects of grain sizes, shapes, and distribution on minimum sizes of representative volume elements of cubic polycrystals, Mech. Mater. 36 (12) (2004) 1217–1229, http://dx.doi.org/10.1016/j.mechmat.2003.11.002.
[60] S.I. Ranganathan, M. Ostoja-Starzewski, Scaling function, anisotropy and the size of RVE in elastic random polycrystals, J. Mech. Phys. Solids 56 (9) (2008) 2773–2791, http://dx.doi.org/10.1016/j.jmps.2008.05.001.
[61] A. Alqaraawi, M. Schuessler, P. Weiß, E. Costanza, N. Berthouze, Evaluating saliency map explanations for convolutional neural networks: a user study, in: IUI '20: Proc. 25th I. Conf. Intelligent User Interfaces, Cagliari, Italy, 2020, pp. 275–285, http://dx.doi.org/10.1145/3377325.3377519.
[62] J. Adebayo, J. Gilmer, M. Muelly, I. Goodfellow, M. Hardt, B. Kim, Sanity checks for saliency maps, in: S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), Adv. Neural Inf. Process. Syst., Vol. 31, Curran Associates, Inc. 2018.
[63] A. Atrey, K. Clary, D. Jensen, Exploratory not explanatory: Counterfactual analysis of saliency maps for deep reinforcement learning, 2020, arXiv preprint arXiv:1912.05743. http://dx.doi.org/10.48550/arXiv.1912.05743.
[64] G. Yona, D. Greenfeld, Revisiting sanity checks for saliency maps, 2021, arXiv preprint arXiv:2110.14297. http://dx.doi.org/10.48550/arXiv.2110.14297.